

Multi-fidelity Gaussian Process Bandit Optimisation

Kirthevasan Kandasamy
Gautam Dasarathy
Junier Oliva
Jeff Schneider
Barnabás Póczos

Carnegie Mellon University, Pittsburgh, PA, USA

KANDASAMY@CS.CMU.EDU
GAUTAMD@CS.CMU.EDU
JOLIVA@CS.CMU.EDU
SCHNEIDE@CS.CMU.EDU
BAPOCZOS@CS.CMU.EDU

Abstract

In many scientific and engineering applications, we are tasked with the optimisation of an expensive to evaluate black box function f . Traditional methods for this problem assume just the availability of this single function. However, in many cases, cheap approximations to f may be obtainable. For example, the expensive real world behaviour of a robot can be approximated by a cheap computer simulation. We can use these approximations to eliminate low function value regions and use the expensive evaluations to f in a small promising region and speedily identify the optimum. We formalise this task as a *multi-fidelity* bandit problem where the target function and its approximations are sampled from a Gaussian process. We develop a method based on upper confidence bound techniques and prove that it exhibits precisely the above behaviour, hence achieving better regret than strategies which ignore multi-fidelity information. Our method outperforms such naive strategies on several synthetic and real experiments.

1. Introduction

In stochastic bandit optimisation, we wish to optimise a *payoff* function $f : \mathcal{X} \rightarrow \mathbb{R}$ by sequentially querying it and obtaining *bandit feedback*, i.e. when we query at any $x \in \mathcal{X}$, we observe a possibly noisy evaluation of $f(x)$. f is typically expensive and the goal is to identify its maximum while keeping the number of queries as low as possible. Some applications are hyper-parameter tuning in expensive machine learning algorithms (Bergstra et al., 2011; Hutter et al., 2011; Snoek et al., 2012), optimal policy search in complex systems (Hornby et al., 2006; Lizotte et al., 2007; Martinez-Cantin et al., 2007), and scientific experiments (Gonzalez et al., 2014; Parkinson et al., 2006). Historically, bandit problems were studied in set-




Figure 1. Average 5-fold CV log likelihood on datasets of size 300, 3000 on a synthetic KDE task. The crosses are the maxima.

tings where the goal is to maximise the cumulative reward of all queries to the payoff instead of just finding the maximum. Applications in this setting include clinical trials and online advertising (Chakrabarti et al., 2008).

Conventional methods in these settings assume access to only this single expensive function of interest f . We will collectively refer to them as single fidelity methods. In many practical problems however, cheap approximations to f might be available. For instance, when tuning hyper-parameters of learning algorithms, the goal is to maximise a cross validation (CV) score on a training set, which can be expensive if the training set is large. But CV curves tend to vary smoothly with training set size; therefore, we can train and cross validate on small subsets to approximate the CV accuracies of the entire dataset. For a concrete example, consider kernel density estimation (KDE), where we need to tune the bandwidth h of a kernel. Figure 1 shows the CV likelihood against h for a dataset of size $n = 3000$ and a smaller subset of size $n = 300$. The two maximisers are different which is to be expected since optimal hyper-parameters are functions of the training set size. That said, the curve for $n = 300$ approximates the $n = 3000$ curve quite well. Since training/CV on small n is cheap, we can use it to eliminate bad values of the hyper-parameters and reserve the expensive experiments with the entire dataset for the promising candidates (e.g. boxed region in Fig. 1).

In online advertising, the goal is to maximise the cumulative number of clicks over a given period. In the conventional bandit treatment, each query to f is the display of

an ad for a specific time, say one hour. However, we may display ads for shorter intervals, say a few minutes, to approximate its hourly performance. The estimate is biased, as displaying an ad for a longer interval changes user behaviour, but will nonetheless be useful in gauging its long run click through rate. In optimal policy search in robotics and automated driving vastly cheaper computer simulations are used to approximate the expensive real world performance of the system. Scientific experiments can be approximated to varying degrees using less expensive data collection, analysis and computational techniques.

In this paper, we study multi-fidelity bandit optimisation, assuming the availability of cheap approximate functions (fidelities) to the payoff f . **Our contributions** are:

1. We present a formalism for multi-fidelity bandit optimisation using Gaussian Process (GP) assumptions on f and its approximations. We develop a novel algorithm, Multi-Fidelity Gaussian Process Upper Confidence Bound (MF-GP-UCB) for this setting.
2. Our theoretical analysis proves that MF-GP-UCB explores the space at lower fidelities and uses the high fidelities in successively smaller regions to zero in on the optimum. As lower fidelity queries are cheaper, MF-GP-UCB has better regret than single fidelity strategies.
3. Empirically, we demonstrate that MF-GP-UCB outperforms single fidelity methods on a series of synthetic examples, three hyper-parameter tuning tasks and one inference problem in Astrophysics. Our matlab implementation and experiments will be made available.

Related Work

Since the seminal work by Robbins (1952), the multi-armed bandit problem has been studied extensively in the K -armed setting (Agrawal, 1995; Lai & Robbins, 1985). Recently, there has been a surge of interest in the optimism under uncertainty principle for K armed bandits, typified by upper confidence bound (UCB) methods (Audibert et al., 2009; Auer, 2003; Bubeck & Cesa-Bianchi, 2012). UCB strategies have also been used in other bandit tasks with linear (Dani et al., 2008) and Gaussian process (de Freitas et al., 2012; Srinivas et al., 2010) payoffs.

There is a plethora of work on single fidelity methods for global optimisation both with noisy and noiseless evaluations. Some examples are branch and bound techniques such as dividing rectangles (DiRect) (Jones et al., 1993), simulated annealing (Kirkpatrick et al., 1983), genetic algorithms (Koza, 1992) and more (Kawaguchi et al., 2015; Munos, 2011; Póczos et al., 2009). A suite of single fidelity methods in the GP framework closely related to our work is Bayesian Optimisation (BO). Some examples of BO techniques are expected improvement (EI) (Mockus, 1994), probability of improvement (PI) (Jones et al., 1998)

Thompson sampling (Thompson, 1933) and predictive entropy search (Hernández-Lobato et al., 2014). Of particular interest to us is the Gaussian process upper confidence bound (GP-UCB) algorithm of Srinivas et al. (2010).

There has been some work on multi-fidelity optimisation with a plurality of them using GP assumptions (Forrester et al., 2007; Huang et al., 2006; Rajnarayan et al., 2008; Xu et al., 2014). However, to our knowledge, all these treatments are rather heuristic in nature and they neither formalise nor analyse any notion of *regret* in the multi-fidelity setting. In contrast, MF-GP-UCB is an intuitive UCB idea with good theoretical properties. Xiong et al. (2013) study 2-fidelity settings where the goal is to predict the outcome of a complex system with the aid of a simpler one. Multi-fidelity methods are used in the robotics community for reinforcement learning tasks by modeling each fidelity as a Markov decision process (Abbeel et al., 2006; Cutler et al., 2014; Kennedy & O’Hagan, 1998). Zhang & Chaudhuri (2015) study active learning with a cheap weak labeler and an expensive strong labeler. The latter sets of works study problems different to optimisation, which is our focus here. Further, none of the papers above are in the bandit setting where there is a price for exploration.

Our work builds on GP-UCB and other UCB strategies, but the multi-fidelity framework ushers in substantially new theoretical and algorithmic challenges. Section 2 presents our formalism including a notion of regret for the multi-fidelity problem and details the problem set up and assumptions. Section 3 presents our algorithm. The theoretical analysis is given in Appendix C of the supplementary material with a synopsis for the 2-fidelity case given in Section 4. Section 6 presents our experiments. Appendix A.1 tabulates the notation used in the manuscript.

2. Preliminaries

The goal is to maximise a noisy black-box function $f : \mathcal{X} \rightarrow \mathbb{R}$. We will focus on settings where \mathcal{X} is either a finite discrete or compact subset of $[0, r]^d$. We can interact with f only by querying at some $x \in \mathcal{X}$ and obtaining a noisy observation $y = f(x) + \epsilon$. Let $x_* \in \operatorname{argmax}_{x \in \mathcal{X}} f(x)$ and $f_* = f(x_*)$. The goal of a bandit strategy is to maximise the sum of rewards $\sum_{t=1}^n f(\mathbf{x}_t)$ or equivalently minimise the cumulative regret $\sum_{t=1}^n f_* - f(\mathbf{x}_t)$ after n queries to f . Here, $\mathbf{x}_t \in \mathcal{X}$ is the queried point at time t . That is, we compete against an oracle which queries at x_* at all t .

Our primary distinction from usual bandit settings is that we have access to $M - 1$ successively accurate approximations $f^{(1)}, f^{(2)}, \dots, f^{(M-1)}$ to the payoff function $f = f^{(M)}$. We will refer to these approximations as fidelities. We encode the fact that lower fidelities approximate the higher fidelity via the assumption, $\|f^{(M)} - f^{(m)}\|_\infty \leq$

$\zeta^{(m)}$ where, $\zeta^{(1)} > \zeta^{(2)} > \dots > \zeta^{(M)} = 0$. Each query at fidelity m expends a cost $\lambda^{(m)}$ of a resource (such as computational effort or advertising time). While the results are generally applicable, they are most interesting in settings where $\lambda^{(1)} \ll \lambda^{(2)} \ll \dots \ll \lambda^{(M)}$.

A strategy for multi-fidelity bandits is a sequence of query-fidelity pairs $\{(\mathbf{x}_t, \mathbf{m}_t)\}_{t \geq 0}$, where $(\mathbf{x}_n, \mathbf{m}_n)$ could depend on the previous query-observation-fidelity tuples $\{(\mathbf{x}_t, \mathbf{y}_t, \mathbf{m}_t)\}_{t=1}^{n-1}$. Here $\mathbf{y}_t = f^{(\mathbf{m}_t)}(\mathbf{x}_t) + \epsilon$. After n steps we will have queried multiple times at any of the M fidelities. $T_n^{(m)}(x)$ denotes the number of queries at $x \in \mathcal{X}$ at fidelity m after n steps. $T_n^{(m)}(A)$ denotes the same for a subset $A \subset \mathcal{X}$. $\mathcal{D}_n^{(m)} = \{(\mathbf{x}_t, \mathbf{y}_t)\}_{t: \mathbf{m}_t = m}$ denotes the set of query-value pairs at the m^{th} fidelity until time n .

Some smoothness assumptions on $f^{(m)}$'s are needed to make the problem tractable. A standard in Bayesian non-parametric literature is to use a Gaussian process (GP) prior (Rasmussen & Williams, 2006) with covariance kernel κ . In this work we focus on the squared exponential (SE) $\kappa_{\sigma, h}$ and the Matérn $\kappa_{\nu, h}$ kernels as their spectral properties are known and since they are popularly used in practice. Writing $r = \|x - x'\|_2$, they are defined as $\kappa_{\sigma, h}(x, x') = \sigma \exp(-r^2/(2h^2))$, $\kappa_{\nu, h}(x, x') = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}r}{h}\right)^\nu B_\nu\left(\frac{\sqrt{2\nu}r}{h}\right)$, where Γ, B_ν are the Gamma and modified Bessel functions. A convenience the GP framework offers is that posterior distributions are analytically tractable. If $f \sim \mathcal{GP}(0, \kappa)$, and we have observations $\mathcal{D}_n = \{(x_i, y_i)\}_{i=1}^n$, where $y_i = f(x_i) + \epsilon$ and $\epsilon \sim \mathcal{N}(0, \eta^2)$ is Gaussian noise, the posterior distribution for $f(x)|\mathcal{D}_n$ is also Gaussian $\mathcal{N}(\mu_n(x), \sigma_n^2(x)^2)$ with,

$$\mu_n(x) = \mathbf{k}^\top \Delta^{-1} Y, \quad \sigma_n^2(x) = \kappa(x, x) - \mathbf{k}^\top \Delta^{-1} \mathbf{k}. \quad (1)$$

Here, $Y \in \mathbb{R}^n$ with $Y_i = y_i$, $\mathbf{k} \in \mathbb{R}^n$ with $\mathbf{k}_i = \kappa(x, x_i)$ and $\Delta = \mathbf{K} + \eta^2 I \in \mathbb{R}^{n \times n}$ where $\mathbf{K}_{i,j} = \kappa(x_i, x_j)$.

In keeping with the above, we make the following assumptions on our problem.

- A1. $f^{(m)} \sim \mathcal{GP}(0, \kappa)$ for all $m = 1, \dots, M$.
- A2. $\|f^{(M)} - f^{(m)}\|_\infty \leq \zeta^{(m)}$ for all $m = 1, \dots, M$.
- A3. $\|f^{(M)}\|_\infty \leq B$.

The purpose of A3 is primarily to define the regret. In Remark 7, Appendix A.3 we argue that these assumptions are probabilistically valid, i.e. the latter two events occur with nontrivial probability when we sample the $f^{(m)}$'s from a GP. So a generative mechanism would keep sampling the functions and deliver them when the conditions hold true. All probabilities henceforth, are conditioned on the above constraints being satisfied. A point $x \in \mathcal{X}$ can be queried at any of the M fidelities. When we query at fidelity m , we observe $y = f^{(m)}(x) + \epsilon$ where $\epsilon \sim \mathcal{N}(0, \eta^2)$.

We now present our notion of cumulative regret $R(\Lambda)$ after spending capital Λ of a resource in the multi-fidelity setting. $R(\Lambda)$ should reduce to the conventional definition of regret for any single fidelity strategy that queries only at M^{th} fidelity. As only the optimum of $f = f^{(M)}$ is of interest to us, queries at fidelities less than M should yield the lowest possible reward, $(-B)$ according to A3. We define the instantaneous reward q_t to be,

$$q_t = \begin{cases} -B & \text{if } \mathbf{m}_t \neq M, \\ f^{(M)}(\mathbf{x}_t) & \text{if } \mathbf{m}_t = M. \end{cases} \quad (2)$$

Also, denote the instantaneous regret by $r_t = f_\star - q_t$. $R(\Lambda)$ should also factor in the costs of the fidelity of each query. Finally, we should also receive $(-B)$ reward for any unused capital. Accordingly, we define $R(\Lambda)$ as,

$$\begin{aligned} R(\Lambda) &= \Lambda f_\star - \left[\sum_{t=1}^N \lambda^{(m_t)} q_t + \left(\Lambda - \sum_{t=1}^N \lambda^{(m_t)} \right) (-B) \right] \\ &= \left(\Lambda - \sum_{t=1}^N \lambda^{(m_t)} \right) (f_\star + B) + \sum_{t=1}^N \lambda^{(m_t)} r_t. \end{aligned} \quad (3)$$

Here, N is the (random) number of queries at all fidelities within capital Λ , i.e. the largest n such that $\sum_{t=1}^n \lambda^{(m_t)} \leq \Lambda$. Lower fidelity queries use up our capital but yield the lowest possible reward. The goal however, is to leverage information from the cheap lower fidelities to query prudently at the higher fidelities and hence obtain better regret.

According to (3) above, we wish to compete against an oracle that uses all its capital Λ to query x_\star at the M^{th} fidelity. $R(\Lambda)$ is at best 0 when we follow the oracle and at most $2\Lambda B$. Our goal is a strategy that has small regret for all values of (sufficiently large) Λ , i.e. the equivalent of an anytime strategy (as opposed to a fixed time horizon strategy) in the usual bandit setting. For the purpose of optimisation, we define the *simple regret* as $S(\Lambda) = \min_t r_t = f_\star - \max_t q_t$. $S(\Lambda)$ is the difference between f_\star and the best highest fidelity query (and $f_\star + B$ if we have never queried at fidelity M). Since $S(\Lambda) \leq \frac{1}{\Lambda} R(\Lambda)$, any strategy with asymptotic sublinear regret $\lim_{\Lambda \rightarrow \infty} \frac{1}{\Lambda} R(\Lambda) = 0$, is also a consistent procedure for optimisation.

A note on UCB strategies for bandits: Sequential optimisation methods adopting UCB principles maintain a high probability upper bound $\varphi_t : \mathcal{X} \rightarrow \mathbb{R}$ for $f(x)$ for all $x \in \mathcal{X}$. φ_t typically takes the form $\varphi_t(\cdot) = \widehat{\mathcal{M}}(\cdot) + \widehat{\mathcal{U}}(\cdot)$ where $\widehat{\mathcal{M}}, \widehat{\mathcal{U}}$ are the estimated mean and a confidence band from the previous $t - 1$ queries. At time t , we query at the maximiser $\mathbf{x}_t = \operatorname{argmax}_{x \in \mathcal{X}} \varphi_t(x)$ of the UCB. The key intuition in UCB strategies is that $\widehat{\mathcal{M}}$ encourages an exploitative strategy—in that we want to query where we know the function is high—and $\widehat{\mathcal{U}}$ encourages an explorative strategy—in that we want to query at regions we




Figure 2. Illustration of GP-UCB. The solid black line is $f(x)$ and the dashed blue line is $\varphi_t(x)$. The observations until $t-1$ are shown as black crosses. At time t , we query at the maximiser $\mathbf{x}_t = \operatorname{argmax}_{x \in \mathcal{X}} \varphi_t(x)$ shown via the red star.

are uncertain about f lest we miss out on high valued regions. For example, in the GP framework, GP-UCB uses $\varphi_t(x) = \mu_{t-1}(x) + \beta_t^{1/2} \sigma_{t-1}(x)$ where μ_{t-1}, σ_{t-1} are the posterior mean and standard deviation of the GP conditioned on the previous $t-1$ queries. μ_{t-1} is the estimated mean explained in the previous paragraph, and for appropriately chosen $\beta_t, \beta_t^{1/2} \sigma_{t-1}^{(m)}$ forms a confidence band. We have illustrated GP-UCB in Figure 2 and reviewed the algorithm along with its theoretical properties in Appendix A.2.

3. MF-GP-UCB

The proposed algorithm, MF-GP-UCB, will also maintain a UCB for $f^{(M)}$ obtained via the previous queries at all fidelities. Denote the mean and standard deviation of $f^{(m)}$ conditioned on $\mathcal{D}_t^{(m)}$ to be $\mu_t^{(m)}(x) = \mathbb{E}[f^{(m)}(x) | \mathcal{D}_t^{(m)}]$, $\sigma_t^{(m)}(x) = (\mathbb{V}[f^{(m)}(x) | \mathcal{D}_t^{(m)}])^{1/2}$. Then define,

$$\begin{aligned} \varphi_t^{(m)}(x) &= \mu_{t-1}^{(m)}(x) + \beta_t^{1/2} \sigma_{t-1}^{(m)}(x) + \zeta^{(m)} \\ \varphi_t(x) &= \min_{m=1, \dots, M} \varphi_t^{(m)}(x) \end{aligned} \quad (4)$$

For appropriately chosen $\beta_t, \mu_{t-1}^{(m)}(x) + \beta_t^{1/2} \sigma_{t-1}^{(m)}(x)$ will upper bound $f^{(m)}(x)$. By A2, $\varphi_t^{(m)}(x)$ upper bounds $f^{(M)}(x)$ for all m . We have M such bounds, and their minimum $\varphi_t(x)$ gives the best bound. Our next query is at the maximiser of this UCB, $\mathbf{x}_t = \operatorname{argmax}_{x \in \mathcal{X}} \varphi_t(x)$.

Next we need to determine which fidelity to query at. Consider any $m < M$. The $\zeta^{(m)}$ conditions on $f^{(m)}$ constrain the value of $f^{(M)}$ – the confidence band $\beta_t^{1/2} \sigma_{t-1}^{(m)}$ in the UCB for $f^{(m)}$ is lengthened by $\zeta^{(m)}$ to obtain confidence on $f^{(M)}$. If $\beta_t^{1/2} \sigma_{t-1}^{(m)}(\mathbf{x}_t)$ for $f^{(m)}$ is large, it means that we haven’t constrained $f^{(m)}$ sufficiently well at \mathbf{x}_t and should query at the m^{th} fidelity. On the other hand, querying indefinitely in the same region to reduce $\beta_t^{1/2} \sigma_{t-1}^{(m)}$ in that region will not help us much as the $\zeta^{(m)}$ elongation caps off how much we can learn about $f^{(M)}$ from $f^{(m)}$; i.e. even if we knew $f^{(m)}$ perfectly, we will only have constrained $f^{(M)}$ to within a $\pm \zeta^{(m)}$ band. Our algorithm captures this simple intuition. Having selected \mathbf{x}_t , we begin by checking

at the first fidelity. If $\beta_t^{1/2} \sigma_{t-1}^{(1)}(\mathbf{x}_t)$ is smaller than a threshold γ , we proceed to check the second fidelity and continue in a similar fashion. If at any stage $\beta_t^{1/2} \sigma_{t-1}^{(m)}(\mathbf{x}_t) > \gamma$ we query at fidelity $\mathbf{m}_t = m$. If we proceed all the way to fidelity M , we query at $\mathbf{m}_t = M$. We summarise the resulting procedure in Algorithm 1.

Algorithm 1 MF-GP-UCB

Input: kernel κ , bounds $(\zeta^{(m)})_{m=1}^M$, parameter γ .

- For $m = 1, \dots, M$:
 $\mathcal{D}_0^{(m)} \leftarrow \emptyset, (\mu_0^{(m)}, \sigma_0^{(m)}) \leftarrow (\mathbf{0}, \kappa^{1/2})$.
 - for $t = 1, 2, \dots$
 1. $\mathbf{x}_t \leftarrow \operatorname{argmax}_{x \in \mathcal{X}} \varphi_t(x)$. (See Equation (4))
 2. **for** $m = 1, \dots, M$:
 $\text{if } \beta_t^{1/2} \sigma_{t-1}^{(m)}(\mathbf{x}_t) > \gamma, \text{ break};$
 $\mathbf{m}_t = m$.
 3. $\mathbf{y}_t \leftarrow \text{Query } f^{(\mathbf{m}_t)} \text{ at } \mathbf{x}_t$.
 4. $\mathcal{D}_t^{(\mathbf{m}_t)} \leftarrow \mathcal{D}_{t-1}^{(\mathbf{m}_t)} \cup \{(\mathbf{x}_t, \mathbf{y}_t)\}$.
 5. $\mathcal{D}_t^{(m)} \leftarrow \mathcal{D}_{t-1}^{(m)}$ for all $m \neq \mathbf{m}_t$.
 6. Obtain $\mu_t^{(\mathbf{m}_t)}, \sigma_t^{(\mathbf{m}_t)}$ conditioned on $\mathcal{D}_t^{(\mathbf{m}_t)}$.
-

Figure 3 illustrates MF-GP-UCB via a 2-fidelity simulation. At the initial stages, MF-GP-UCB is mostly exploring \mathcal{X} in the first fidelity. $\beta_t^{1/2} \sigma_{t-1}^{(1)}$ is large and we are yet to constrain $f^{(1)}$ well to proceed to $m = 2$. By $t = 14$, we have constrained $f^{(1)}$ well around the optimum and therefore have started querying at the second fidelity in this region. Notice that $\varphi_t^{(2)}$ dips to change φ_t in that region. MF-GP-UCB has identified the maximum x_* with just 3 queries to $f^{(2)}$. In Appendix B we provide an extended version of this simulation and discuss further insights.

Finally, we make an essential observation. The posterior distribution of any $f^{(m)}(x)$ conditioned on previous queries at all fidelities $\mathcal{D}_t^{(m)}, m = 1, \dots, M$ is not Gaussian due to the $\zeta^{(m)}$ constraints (A2). However, $|f^{(m)}(x) - \mu_{t-1}^{(m)}(x)| < \beta_t^{1/2} \sigma_{t-1}^{(m)}(x)$ holds with high probability, since, by conditioning only on $\mathcal{D}_t^{(m)}$ we have Gaussianity for $f^{(m)}(x)$ (See Lemma 11, Appendix C.1). Next we summarise our main theoretical contributions.

4. Summary of Theoretical Results

For purely pedagogical reasons we present our results for the $M = 2$ case here. Appendix C contains theorem statements and proofs for general M . We also ignore constants and polylog terms whenever they are subdominant. \lesssim, \asymp denote inequality and equality ignoring constants. We begin by defining the *Maximum Information Gain* (MIG) which characterises the statistical difficulty of GP bandits.

Definition 1. (Maximum Information Gain) Let $f \sim$




Figure 3. Illustration of MF-GP-UCB for a 2-fidelity problem which has been initialised with 5 random points at the first fidelity. In the top figures, the solid lines in brown and blue are $f^{(1)}, f^{(2)}$ respectively, and the dashed lines are $\varphi_t^{(1)}, \varphi_t^{(2)}$. The solid green line is $\varphi_t = \min(\varphi_t^{(1)}, \varphi_t^{(2)})$. The small crosses are queries from 1 to $t-1$ and the red star is the maximiser of φ_t , i.e. the next query \mathbf{x}_t . x_* , the optimum of $f^{(2)}$ is shown in magenta. In the bottom figures, the solid orange line is $\beta_t^{1/2} \sigma_{t-1}^{(1)}$ and the dashed black line is γ . When $\beta_t^{1/2} \sigma_{t-1}^{(1)}(\mathbf{x}_t) \leq \gamma$ we play at fidelity $\mathbf{m}_t = 2$ and otherwise at $\mathbf{m}_t = 1$. See Fig. 7 in Appendix B for an extended simulation.

$\mathcal{GP}(\mathbf{0}, \kappa)$. Consider any $A \subset \mathbb{R}^d$ and let $\tilde{A} = \{x_1, \dots, x_n\} \subset A$ be a finite subset. Let $f_{\tilde{A}}, \epsilon_{\tilde{A}} \in \mathbb{R}^n$ such that $(f_{\tilde{A}})_i = f(x_i)$ and $(\epsilon_{\tilde{A}})_i \sim \mathcal{N}(0, \eta^2)$ and $y_{\tilde{A}} = f_{\tilde{A}} + \epsilon_{\tilde{A}}$. Let I denote the Shannon Mutual Information. The Maximum Information Gain of A is

$$\Psi_n(A) = \max_{\tilde{A} \subset A, |\tilde{A}|=n} I(y_{\tilde{A}}; f_{\tilde{A}}).$$

The MIG, which depends on the kernel and the set A , will be an important quantity in our analysis. For a given kernel it typically scales with the volume of A . For e.g. if $A = [0, r]^d$ then $\Psi_n(A) \in \mathcal{O}(r^d \Psi_n([0, 1]^d))$. For the SE kernel, $\Psi_n([0, 1]^d) \in \mathcal{O}((\log(n))^{d+1})$ and for the Matérn kernel, $\Psi_n([0, 1]^d) \in \mathcal{O}(n^{\frac{d(d+1)}{2v+d(d+1)}} \log(n))$ (Srinivas et al., 2010).

Fundamental to the 2-fidelity problem are the following sets: the “good” set $\mathcal{X}_g = \{x \in \mathcal{X}; f_* - f^{(1)}(x) \leq \zeta^{(1)}\}$ and the “bad” set $\mathcal{X}_b = \mathcal{X} \setminus \mathcal{X}_g$. \mathcal{X}_g is good because it is a high valued region for $f^{(2)}(x)$. For instance, for all $x \in \mathcal{X}_g$, $f^{(2)}(x)$ is at most $2\zeta^{(1)}$ away from the optimum. More interestingly, when $\zeta^{(1)}$ is small, i.e. when $f^{(1)}$ is a good approximation to $f^{(2)}$, \mathcal{X}_g will be much smaller than \mathcal{X} . This is precisely the target domain for this research. For example, in optimal policy search in robotics a cheap computer simulator can be used to eliminate several bad policies and we could reserve the real world experiments for the promising candidates. If a multi-fidelity strategy were to use the second fidelity queries only in \mathcal{X}_g , then the regret will only have $\Psi_n(\mathcal{X}_g)$ dependence after n high fidelity queries. In contrast, a strategy that only operates at the highest fidelity (e.g. GP-UCB) will have $\Psi_n(\mathcal{X})$ dependence. In the scenario described above $\Psi_n(\mathcal{X}_g) \ll \Psi_n(\mathcal{X})$, and the multi-fidelity strategy will have significantly better regret than a single fidelity strategy. We will see that MF-GP-UCB roughly achieves this goal.

Recall, N is the (random) number of queries by a multi-fidelity strategy within capital Λ at either fidelity, $N = T_N^{(1)}(\mathcal{X}) + T_N^{(2)}(\mathcal{X})$. Let $n_\Lambda = \lfloor \Lambda / \lambda^{(2)} \rfloor$ be the (non-random) number of queries by a single fidelity method operating only at the second fidelity. As $\lambda^{(1)} \ll \lambda^{(2)}$, N could be large for an arbitrary multi-fidelity method. However, our analysis reveals that for MF-GP-UCB, N is on the order of n_Λ and most 2nd fidelity queries are roughly in \mathcal{X}_g .

We first present our theorem for $M = 2$ for finite discrete domains. Define, the following set parametrised by ρ , $\tilde{\mathcal{X}}_{g,\rho} = \{x; f_* - f^{(1)}(x) \leq \zeta^{(1)} + \rho\gamma\}$. MF-GP-UCB will query at the second fidelity mostly in $\tilde{\mathcal{X}}_{g,\rho}$; the choice of γ , results in a slight inflation of \mathcal{X}_g . As $\gamma \rightarrow 0$, $\tilde{\mathcal{X}}_{g,\rho} \rightarrow \mathcal{X}_g$. In addition, define $\tilde{\mathcal{X}}_{b,\rho} = \mathcal{X} \setminus \tilde{\mathcal{X}}_{g,\rho}$.

Theorem 2. Let $\mathcal{X} \subset [0, r]^d$ be finite and $f^{(1)}, f^{(2)} \sim \mathcal{GP}(\mathbf{0}, \kappa)$ satisfy Assumptions A2, A3. Pick $\delta \in (0, 1)$ and run MF-GP-UCB with $\beta_t \asymp \log(|\mathcal{X}|t/\delta)$. Then, with probability $> 1 - \delta$, for all sufficiently large Λ and for all $\alpha \in (0, 1)$, there exists ρ depending on α such that,

$$R(\Lambda) \lesssim \lambda^{(1)} \xi |\tilde{\mathcal{X}}_{b,\rho}| \beta_{n_\Lambda} + \lambda^{(1)} \frac{|\tilde{\mathcal{X}}_{g,\rho}| \beta_{n_\Lambda}}{\gamma^2} + \lambda^{(2)} \sqrt{\beta_{n_\Lambda} |\mathcal{X}| n_\Lambda^\alpha \Psi_{|\mathcal{X}|n_\Lambda^\alpha}(\mathcal{X})} + \lambda^{(2)} \sqrt{\beta_{n_\Lambda} n_\Lambda \Psi_{n_\Lambda}(\tilde{\mathcal{X}}_{g,\rho})}.$$

Here ξ is a problem dependent constant determined by the values $f^{(m)}(x)$ for all m, x .

It is instructive to compare the above rates against that for GP-UCB (see Theorem 4, Appendix A.2). By dropping the common and subdominant terms, the rate for MF-GP-UCB is $\lambda^{(2)} \sqrt{n_\Lambda \Psi_{n_\Lambda}(\tilde{\mathcal{X}}_{g,\rho})} + \lambda^{(2)} \sqrt{n_\Lambda^\alpha \Psi_{n_\Lambda^\alpha}(\mathcal{X})}$ whereas for GP-UCB it is $\lambda^{(2)} \sqrt{n_\Lambda \Psi_{n_\Lambda}(\mathcal{X})}$. For GP-UCB, the MIG for the entire \mathcal{X} is coupled with n_Λ , where as for MF-GP-UCB it is only coupled with the small n_Λ^α term.

MF-GP-UCB also has an MIG coupling with n_Λ , but this is on the set $\tilde{\mathcal{X}}_{g,\rho}$. Let us say that $\tilde{\mathcal{X}}_{g,\rho}$ is contained in a (possibly disconnected) region of volume $\nu \ll r^d$, then $\Psi_{n_\Lambda}(\tilde{\mathcal{X}}_{g,\rho}) \ll \Psi_{n_\Lambda}(\mathcal{X})$ and MF-GP-UCB has better regret than GP-UCB. This directly captures previously mentioned intuitions: use lower fidelities to eliminate bad regions of \mathcal{X} and query the higher fidelities only at promising regions. The price we have to pay is that part of our capital is spent on first fidelity queries. However these terms are logarithmic in n_Λ and more importantly the queries are cheap ($\lambda^{(1)} \ll \lambda^{(2)}$). In M fidelities, there will be a hierarchy of $M - 1$ successively smaller “good sets” and MF-GP-UCB will query at fidelity m only in the corresponding set.

The bound also illustrates the trade-off on the choice of γ . As $\gamma \rightarrow 0$, $1/\gamma^2$ increases but the set $\tilde{\mathcal{X}}_{g,\rho}$ decreases. This observation however, is of little practical use as there are problem dependent terms. In Section 5 we describe some heuristics to set γ which worked well in our experiments.

Proof sketch of Theorem 2: We control the number of first fidelity queries in $\tilde{\mathcal{X}}_{g,\rho}$, $\tilde{\mathcal{X}}_{b,\rho}$ and second fidelity queries in $\tilde{\mathcal{X}}_{b,\rho}$. For $x \in \tilde{\mathcal{X}}_{b,\rho}$, we show $T_n^{(1)}(x) \lesssim \log(n)$ by arguing that it won't be the UCB maximiser too many times. For $x \in \tilde{\mathcal{X}}_{g,\rho}$, $T_n^{(1)}(x) \lesssim \log(n)/\gamma^2$ using the switching criterion (step 2) in Algorithm 1. For $x \in \tilde{\mathcal{X}}_{b,\rho}$, we show $T_n^{(2)}(x) \lesssim n^\alpha$ by appealing to previous first fidelity queries. Therefore $n - T_n^{(2)}(\tilde{\mathcal{X}}_{g,\rho}) \in o(n)$, i.e. most queries are at the second fidelity and confined to $\tilde{\mathcal{X}}_{g,\rho}$. Finally, We invoke techniques from Srinivas et al. (2010) to control the regret. However, unlike them we analyse the MIG of $\tilde{\mathcal{X}}_{g,\rho}$ and $\tilde{\mathcal{X}}_{b,\rho}$ separately using the bound on $T_n^{(1)}(\tilde{\mathcal{X}}_{b,\rho})$ and obtain a tighter bound on $R(\Lambda)$. The theorem and proof for general M is in Appendix C.1. ■

We now analyse compact and convex \mathcal{X} . We consider a slightly different inflation $\tilde{\mathcal{X}}_{g,\rho,\tau} = \{x \in \mathcal{X}; f_\star - f^{(1)}(x) \leq \zeta^{(1)} + \max(\tau, \rho\gamma)\}$, of \mathcal{X}_g . Denote $\tilde{\mathcal{X}}_{b,\tau} = \mathcal{X} \setminus \tilde{\mathcal{X}}_{g,\tau}$. In addition, for any given $\alpha > 0$ we will let $\tilde{\mathcal{X}}_{g,\rho,\tau,n} = \{x \in \mathcal{X} : B_2(x, r\sqrt{d}/n^{\alpha/2d}) \cap \tilde{\mathcal{X}}_{g,\rho,\tau} \neq \emptyset\}$ denote a further inflation of $\tilde{\mathcal{X}}_{g,\rho,\tau}$. Here $B_2(x, \epsilon)$ is the L_2 ball of radius ϵ centred at x . Observe that $\forall \alpha > 0$, as $n \rightarrow \infty$, $\tilde{\mathcal{X}}_{g,\rho,\tau,n} \rightarrow \tilde{\mathcal{X}}_{g,\rho,\tau}$. Finally, let $\Omega_\varepsilon(A)$ be the ε -covering number of a set $A \subset \mathcal{X}$ in the L_2 metric.

Theorem 3. Let $\mathcal{X} \subset [0, r]^d$ be compact and convex and $f^{(1)}, f^{(2)} \sim \mathcal{GP}(\mathbf{0}, \kappa)$ satisfy A2, A3. Pick $\delta \in (0, 1)$ and run MF-GP-UCB with $\beta_t \asymp d \log(t/\delta)$. Then, with probability $> 1 - \delta$, for sufficiently large Λ and for all $\alpha \in (0, 1)$, and $\tau > 0$, there exists ρ depending on α such that,

$$R(\Lambda) \lesssim \lambda^{(1)} \sqrt{n_\Lambda \beta_{n_\Lambda} \Psi_{n_\Lambda}(\mathcal{X})} + \lambda^{(1)} \frac{\Omega_{\varepsilon_n}(\tilde{\mathcal{X}}_{g,\rho,\tau}) \beta_{n_\Lambda}}{\gamma^2} \\ + \lambda^{(2)} \sqrt{M n_\Lambda^\alpha \beta_{n_\Lambda} \Psi_{M n_\Lambda^\alpha}(\mathcal{X})} + \lambda^{(2)} \sqrt{n_\Lambda \beta_{n_\Lambda} \Psi_{n_\Lambda}(\tilde{\mathcal{X}}_{g,\rho,\tau,n_\Lambda})} \\ \text{Here } \varepsilon_n \text{ is } \asymp \frac{\gamma}{\sqrt{\beta_n}} \text{ for the SE kernel and } \asymp \frac{\gamma^2}{\beta_n} \text{ for Matérn.}$$

There are a number differences from the discrete case. Previously, the first term had only $\log(n_\Lambda)$ dependence whereas here we have $\lambda^{(1)} \sqrt{n_\Lambda \Psi_{n_\Lambda}(\mathcal{X})}$. This is insignificant compared to GP-UCB when $\lambda^{(1)} \ll \lambda^{(2)}$. The second term is $\mathcal{O}(r'^d (\log(n))^{1+d/2} / \gamma^{d+2})$ for the SE kernel and $\mathcal{O}(r'^d (\log(n))^{1+d} / \gamma^{2d+2})$ for the Matérn kernel, where $r' = \text{diam}(\tilde{\mathcal{X}}_{g,\rho,\tau})$. Finally, the last term has $\tilde{\mathcal{X}}_{g,\rho,\tau,n_\Lambda}$ instead of $\tilde{\mathcal{X}}_{g,\rho,\tau}$. However, as $\tilde{\mathcal{X}}_{g,\rho,\tau,n_\Lambda}$ shrinks to $\tilde{\mathcal{X}}_{g,\rho,\tau}$ polynomially fast, whenever $\text{vol}(\tilde{\mathcal{X}}_{g,\rho,\tau}) \ll \text{vol}(\mathcal{X})$, MF-GP-UCB achieves better regret than GP-UCB.

Proof sketch of Theorem 3: We bound the number of first fidelity plays in $\tilde{\mathcal{X}}_{g,\rho,\tau}$, $T_n^{(1)}(\tilde{\mathcal{X}}_{g,\rho,\tau})$ via a covering argument. For this, we establish an upper bound on the posterior variance of a GP in a subset $A \subset \mathcal{X}$ after a certain number of queries in A . To bound $T_n^{(1)}(\tilde{\mathcal{X}}_{b,\rho,\tau})$, we consider only the first fidelity queries in an modified game at $f^{(1)}$ and show that MF-GP-UCB achieves good regret in this game. The purpose of the τ term in the inflation is to create a separation between the two sets – in the discrete case, the separation is guaranteed due to the discrete $f^{(1)}$ values and is factored in the ξ term. We show $T_n^{(2)}(\tilde{\mathcal{X}}_{b,\rho,\tau}) \lesssim n^\alpha$ by appealing to previous first fidelity queries. The remainder of the analysis follows similar to the discrete case by separately analysing the MIGs of $\tilde{\mathcal{X}}_{g,\rho,\tau,n}$ and $\tilde{\mathcal{X}}_{b,\rho,\tau}$. The theorem and proof for general M is in Appendix C.2. ■

Theorems 2, 3 can be generalised to cases where each $f^{(m)}$ is sampled from a different kernel $\kappa^{(m)}$, and observation noises $\eta^{(m)}$ are different. In fact, our practical implementation uses different kernels. A polylog(n_Λ) rate is also possible for the discrete case, similar to Theorem 1 in Dani et al. (2008), but it will have worse dependence on Ψ_{n_Λ} . We stick to the above form which is comparable to the results of Srinivas et al. (2010). We believe that the $\zeta^{(m)}$ constraint can be relaxed to hold only in some neighborhood of x_\star for MF-GP-UCB to still have good regret; i.e. $f^{(M)}(x) - f^{(m)}(x) \leq \zeta^{(m)}$ for all x in a neighborhood $\nu^{(m)}$ of x_\star . Our proof will need some massaging. Just like any nonparametric method, our algorithm has exponential dependence on dimension. By assuming additional structure in the bandit problem (Djolonga et al., 2013; Kandamay et al., 2015) we may alleviate this problem. Finally, we remind the reader that the above rates also translate to bounds on the simple regret $S(\Lambda)$ for optimisation.

5. Implementation Details

Our implementation uses some standard techniques in Bayesian optimisation to learn the GP kernel such as initialisation with random queries and periodic marginal likelihood maximisation. We also use a more aggressive choice for β_t than stipulated by our analysis. The above techniques might be already known to a reader familiar with




Figure 4. The simple regret $S(\Lambda)$ against the spent capital Λ on synthetic functions. The title states the function, its dimensionality, the number of fidelities and the costs we used for each fidelity in the experiment. All curves barring DiRect (which is a deterministic), were produced by averaging over 20 experiments. The error bars indicate one standard error. See Figures 10, 11 12 in Appendix D for more synthetic results. The last panel shows the number of queries at different function values at each fidelity for the Hartmann-3D example.

the BO literature. We have elaborated these in Appendix B but now focus on the γ and $\zeta^{(m)}$ parameters of our method.

Algorithm 1 assumes that the $\zeta^{(m)}$'s are given with the problem description, which is hardly the case in practice. In our implementation, instead of having to deal with $M - 1$, $\zeta^{(m)}$ values we set $(\zeta^{(1)}, \zeta^{(2)}, \dots, \zeta^{(M-1)}) = ((M - 1)\zeta, (M - 2)\zeta, \dots, \zeta)$ so we only have one value ζ . For e.g. this corresponds to $\|f^{(m)} - f^{(m-1)}\|_\infty \leq \zeta$ which is stronger than Assumption A2. Initially, we start with small ζ . Whenever we query at any fidelity $m > 1$ we also check the posterior mean of the $(m - 1)^{\text{th}}$ fidelity. If $|f^{(m)}(\mathbf{x}_t) - \mu_{t-1}^{(m-1)}(\mathbf{x}_t)| > \zeta$, we query again at \mathbf{x}_t , but at the $(m - 1)^{\text{th}}$ fidelity. If $|f^{(m)}(\mathbf{x}_t) - f^{(m-1)}(\mathbf{x}_t)| > \zeta$, we update ζ to twice the violation.

To set γ we use the following intuition: if the algorithm, is stuck at a lower fidelity for too long then γ is probably too small. We start with a small value of γ . If the algorithm does not query above the m^{th} fidelity for more than $\lambda^{(m+1)}/\lambda^{(m)}$ iterations, we increase γ to twice the current value. We found our implementation to be fairly robust even recovering from fairly bad approximations at the lower fidelities (see Appendix D.2 for an experiment).

6. Experiments

We present experiments for compact \mathcal{X} since it is the more practically relevant setting. We compare MF-GP-UCB to the following single fidelity optimisation methods: GP-UCB, EI, PI, RAND and DiRect. RAND is where we query \mathcal{X} uniformly at random and take the maximum. For all GP methods (GP-UCB, EI, PI) we use the SE kernel. We could not find software for any multi-fidelity methods in previous literature. We implemented the GP based method of Huang et al. (2006) to our best interpretation but found that it did not perform as desired – it hardly queried above the first fidelity. Hence, we exclude it from the comparisons.

A straightforward way to incorporate lower fidelity infor-

mation to single fidelity GP methods is to query at lower fidelities and use them in learning the kernel by jointly maximising the GP marginal likelihood. While the idea seems natural, we got mixed results in practice. On some problems this approach improved the performance of all GP methods (including MF-GP-UCB), but on others all methods worked poorly. One explanation is that while lower fidelities approximate function values, they are not always best described by the same kernel. The results presented do not use lower fidelity queries to learn the kernel as it was more robust. For MF-GP-UCB, the kernels for each $f^{(m)}$ were learned independently using queries at that fidelity.

6.1. Synthetic Examples

We use the Currin exponential ($d = 2$), Park ($d = 4$) and Borehole ($d = 8$) functions in $M = 2$ fidelity experiments and the Hartmann functions in $d = 3$ and 6 with $M = 3$ and 4 fidelities respectively. The first three functions are taken from previous multi-fidelity literature (Xiong et al., 2013) while we tweaked the Hartmann functions to obtain the lower fidelities for the latter two cases. We show the simple regret $S(\Lambda)$ against capital Λ for the Borehole and Hartmann-3D functions in Fig. 4 with the rest deferred to Appendix D due to space constraints. Very clearly, MF-GP-UCB outperforms the single fidelity methods. Appendix D also contains results for the cumulative regret $R(\Lambda)$ and the formulae for these functions.

The third panel of Fig. 4 shows a histogram of the number of queries at each fidelity after 184 queries of MF-GP-UCB, for different ranges of $f^{(3)}(x)$ for the Hartmann-3D example. Many of the queries at the low $f^{(3)}$ values are at fidelity 1, but as we progress they decrease and the second fidelity queries increase. The third fidelity dominates very close to the optimum but is used sparingly elsewhere. This corroborates the prediction in our analysis that MF-GP-UCB uses low fidelities to explore and successively higher fidelities at more promising regions to zero in on x_* . (Also see the simulation in Fig. 7, Appendix B.)




Figure 5. Results on the hyper-parameter tuning experiments. The title states the experiment, dimensionality (number of hyperparameters) and training set size at each fidelity. All curves were produced by averaging over 10 experiments. The lengths of the curves are different in time as we ran each method for a pre-specified number of iterations and they concluded at different times.




Figure 6. Results on the supernova inference problem. The y -axis is the log likelihood so higher is better. RAND is not visible as it performed very poorly. See caption in Fig. 5 for more details.

One should read synthetic results with a grain of salt. By making the lower fidelity costs small (large) and their approximations good (bad), the experiment can be designed favourably (unfavourably) for multi-fidelity methods. The litmus test for this framework lies in how well cheap experiments approximate the highest fidelity in practical problems and if a method can exploit this information.

6.2. Real Experiments

We present results on three hyper-parameter tuning tasks (results in Fig. 5), and a maximum likelihood inference task in Astrophysics. We compare methods on computation time since that is the “cost” in all experiments. We include the processing time for each method in the comparison (i.e. the cost of determining the next query).

Classification using SVMs (SVM): We trained an SVM on the magic gamma dataset using the sequential minimal optimisation algorithm to an accuracy of 10^{-12} . The goal is to tune the kernel bandwidth and the soft margin coefficient in the ranges $(10^{-3}, 10^1)$ and $(10^{-1}, 10^5)$ respectively on a dataset of size 2000. We set this up as a $M = 2$ fidelity experiment with the entire training set at the second fidelity and 500 points at the first. Each query was 5-fold cross validation on these training sets.

Regression using additive kernels (AddKRR): We used the Additive kernel ridge regression method of Kandasamy & Yu (2015) on the 4-dimensional coal power plant dataset. We tuned the 6 hyper-parameters –the regularisation penalty, the kernel scale and the kernel bandwidth for each dimension– each in the range $(10^{-3}, 10^4)$ using 5-fold cross validation. This experiment used $M = 3$ and 2000, 4000, 8000 points at each fidelity respectively.

Viola & Jones face detection (V&J): The Viola & Jones cascade face classifier (Viola & Jones, 2001), which uses a cascade of weak classifiers, is a popular method for face detection. To classify an image, we pass it through each classifier. If at any point the classifier score falls below a threshold, the image is classified as negative. If it passes through the cascade, then it is classified as positive. One of the more popular implementations comes with OpenCV and uses a cascade of 22 weak classifiers. The threshold values in the OpenCV implementation are pre-set based on some heuristics and there is no reason to think they are optimal for a given face detection problem. The goal is to tune these 22 thresholds by optimising for them over a training set. We modified the OpenCV implementation to take in the thresholds as parameters. As our domain \mathcal{X} we chose a neighbourhood around the configuration used in OpenCV. We set this up as a $M = 2$ fidelity experiment where the second fidelity used 3000 images from the Viola and Jones face database and the first used just 300. Interestingly, on an independent test set, the configurations found by MF-GP-UCB consistently achieved over 90% accuracy while the OpenCV configuration achieved only 87.4% accuracy.

Type Ia Supernovae: We use Type Ia supernovae data from Davis et al (2007) for maximum likelihood inference on 3 cosmological parameters, the Hubble constant $H_0 \in (60, 80)$, the dark matter fraction $\Omega_M \in (0, 1)$ and the dark energy fraction $\Omega_\Lambda \in (0, 1)$. Unlike typical parametric maximum likelihood problems we see in machine learning, the likelihood is only available as a black-box. It is computed using the Robertson–Walker metric which requires a (one dimensional) numerical integration for each

sample in the dataset. We set this up as a $M = 3$ fidelity task. At the third fidelity, the integration was performed using the trapezoidal rule on a grid of size 10^6 . For the first and second fidelities, we used grids of size 10^2 , 10^4 respectively. The goal is to maximise the likelihood at the third fidelity. The results are given in Fig. 6.

7. Conclusion

We introduced and studied the multi-fidelity bandit problem under Gaussian Process assumptions. Our theorems demonstrate that MF-GP-UCB explores the space using the cheap lower fidelities, and uses the higher fidelity queries on successively smaller regions hence achieving better regret than single fidelity strategies. Experimental results demonstrate the efficacy of our method and more generally, the utility of the multi-fidelity framework.

It will be useful to extend the multi-fidelity framework to other bandit settings and more broadly to other sequential experiment selection problems such as active learning, on-line learning, reinforcement learning etc. This area of research, largely untapped, has potential for great impact on a wide range of real world applications.

Acknowledgements

We wish to thank Bharath Sriperumbudur for the helpful email discussions. This research was partly funded by DOE grant DESC0011114.

References

- Abbeel, Pieter, Quigley, Morgan, and Ng, Andrew Y. Using inaccurate models in reinforcement learning. In *ICML*, 2006.
- Agrawal, Rajeev. Sample Mean Based Index Policies with $O(\log n)$ Regret for the Multi-Armed Bandit Problem. *Advances in Applied Probability*, 1995.
- Audibert, Jean-Yves, Munos, Rémi, and Szepesvári, Csaba. Exploration-exploitation Tradeoff Using Variance Estimates in Multi-armed Bandits. *Theor. Comput. Sci.*, 2009.
- Auer, Peter. Using Confidence Bounds for Exploitation-exploration Trade-offs. *J. Mach. Learn. Res.*, 2003.
- Bergstra, James S., Bardenet, Rémi, Bengio, Yoshua, and Kégl, Balázs. Algorithms for Hyper-Parameter Optimization. In *NIPS*, 2011.
- Brochu, Eric, Cora, Vlad M., and de Freitas, Nando. A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning. *CoRR*, 2010.
- Bubeck, Sébastien and Cesa-Bianchi, Nicolò. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 2012.
- Bull, Adam D. Convergence Rates of Efficient Global Optimization Algorithms. *JMLR*, 2011.
- Chakrabarti, Deepayan, Kumar, Ravi, Radlinski, Filip, and Upfal, Eli. Mortal Multi-Armed Bandits. In *NIPS*, 2008.
- Cutler, Mark, Walsh, Thomas J., and How, Jonathan P. Reinforcement Learning with Multi-Fidelity Simulators. In *ICRA*, 2014.
- Dani, Varsha, Hayes, Thomas P., and Kakade, Sham M. Stochastic Linear Optimization under Bandit Feedback. In *COLT*, 2008.
- Davis et al, T. M. Scrutinizing Exotic Cosmological Models Using ESSENCE Supernova Data Combined with Other Cosmological Probes. *The Astrophysical Journ.*, 2007.
- de Freitas, Nando, Smola, Alex J., and Zoghi, Masrour. Exponential Regret Bounds for Gaussian Process Bandits with Deterministic Observations. In *ICML*, 2012.
- Djolonga, Josip, Krause, Andreas, and Cevher, Volkan. High-Dimensional Gaussian Process Bandits. In *Advances in Neural Information Processing Systems*, 2013.
- Forrester, Alexander I. J., Söbester, András, and Keane, Andy J. Multi-fidelity optimization via surrogate modelling. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 2007.
- Ghosal, Subhashis and Roy, Anindya. Posterior consistency of Gaussian process prior for nonparametric binary regression”. *Annals of Statistics*, 2006.
- Gonzalez, Javier, Longworth, Joseph, James, David, and Lawrence, Neil. Bayesian Optimization for Synthetic Gene Design. In *NIPS Workshop on Bayesian Optimization in Academia and Industry*, 2014.
- Hernández-Lobato, José Miguel, Hoffman, Matthew W, and Ghahramani, Zoubin. Predictive Entropy Search for Efficient Global Optimization of Black-box Functions. In *NIPS*, 2014.
- Hornby, G. S., Globus, A., Linden, D.S., and Lohn, J.D. Automated Antenna Design with Evolutionary Algorithms. *American Institute of Aeronautics and Astronautics*, 2006.
- Huang, D., Allen, T.T., Notz, W.I., and Miller, R.A. Sequential kriging optimization using multiple-fidelity evaluations. *Structural and Multidisciplinary Optimization*, 2006.

- Hutter, Frank, Hoos, Holger H., and Leyton-Brown, Kevin. Sequential Model-based Optimization for General Algorithm Configuration. In *LION*, 2011.
- Jones, D. R., Perttunen, C. D., and Stuckman, B. E. Lipschitzian Optimization Without the Lipschitz Constant. *J. Optim. Theory Appl.*, 1993.
- Jones, Donald R., Schonlau, Matthias, and Welch, William J. Efficient global optimization of expensive black-box functions. *J. of Global Optimization*, 1998.
- Kandasamy, Kirthevasan and Yu, Yaoliang. Additive Approximations in High Dimensional Nonparametric Regression via the SALSA. *CoRR*, abs/1602.00287, 2015.
- Kandasamy, Kirthevasan, Schenider, Jeff, and Póczos, Barnabás. High Dimensional Bayesian Optimisation and Bandits via Additive Models. In *International Conference on Machine Learning*, 2015.
- Kawaguchi, Kenji, Kaelbling, Leslie Pack, and Lozano-Pérez, Tomás. Bayesian Optimization with Exponential Convergence. In *Advances in Neural Information Processing (NIPS)*, 2015.
- Kennedy, M. C. and O’Hagan, A. Predicting the Output from a Complex Computer Code when Fast Approximations are Available. *Biometrika*, 1998.
- Kirkpatrick, S., Gelatt, C. D., and Vecchi, M. P. Optimization by simulated annealing. *SCIENCE*, 1983.
- Koza, John R. *Genetic programming : on the programming of computers by means of natural selection*. Cambridge, Mass. MIT Press, 1992.
- Lai, T. L. and Robbins, Herbert. Asymptotically Efficient Adaptive Allocation Rules. *Advances in Applied Mathematics*, 1985.
- Lizotte, Daniel, Wang, Tao, Bowling, Michael, and Schuurmans, Dale. Automatic gait optimization with gaussian process regression. In *IJCAI*, 2007.
- Martinez-Cantin, R., de Freitas, N., Doucet, A., and Castellanos, J. Active Policy Learning for Robot Planning and Exploration under Uncertainty. In *Proceedings of Robotics: Science and Systems*, 2007.
- Mockus, Jonas. Application of Bayesian approach to numerical methods of global and stochastic optimization. *Journal of Global Optimization*, 1994.
- Munos, R. Optimistic Optimization of Deterministic Functions without the Knowledge of its Smoothness. In *NIPS*, 2011.
- Parkinson, David, Mukherjee, Pia, and Liddle, Andrew R. A Bayesian model selection analysis of WMAP3. *Physical Review*, 2006.
- Póczos, Bárnabás, Abbasi-Yadkori, Yasin, Szepesvri, Csaba, Greiner, Russell, and Sturtevant, Nathan R. Learning when to stop thinking and do something! In *ICML*, 2009.
- Rajnarayan, Dev, Haas, Alex, and Kroo, Ilan. A multi-fidelity gradient-free optimization method and application to aerodynamic design. In *AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 2008.
- Rasmussen, C.E. and Williams, C.K.I. *Gaussian Processes for Machine Learning*. Adaptive computation and machine learning series. University Press Group Limited, 2006.
- Robbins, Herbert. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 1952.
- Snoek, Jasper, Larochelle, Hugo, and Adams, Ryan P. Practical Bayesian Optimization of Machine Learning Algorithms. In *NIPS*, 2012.
- Srinivas, Niranjan, Krause, Andreas, Kakade, Sham, and Seeger, Matthias. Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design. In *International Conference on Machine Learning*, 2010.
- Thompson, W. R. On the Likelihood that one Unknown Probability Exceeds Another in View of the Evidence of Two Samples. *Biometrika*, 1933.
- Viola, Paul A. and Jones, Michael J. Rapid Object Detection using a Boosted Cascade of Simple Features. In *Computer Vision and Pattern Recognition*, 2001.
- Xiong, Shifeng, Qian, Peter Z. G., and Wu, C. F. Jeff. Sequential design and analysis of high-accuracy and low-accuracy computer codes. *Technometrics*, 2013.
- Xu, Jie, Zhang, Si, Huang, Edward, Chen, Chun-Hung, Lee, Loo Hay, and Celik, Nurcin. Efficient Multi-fidelity Simulation Optimization. In *Proceedings of the 2014 Winter Simulation Conference*, 2014.
- Zhang, Chicheng and Chaudhuri, Kamalika. Active Learning from Weak and Strong Labelers. In *NIPS*, 2015.

Appendix

A. Some Ancillary Material

A.1. Table of Notations

M	The number of fidelities.
$f, f^{(m)}$	The payoff function and its m^{th} fidelity approximation. $f^{(M)} = f$.
$\lambda^{(m)}$	The cost for querying at fidelity m .
\mathcal{X}	The domain over which we are optimising f .
x_*, f_*	The optimum point and value of the M^{th} fidelity function.
\bar{A}	The complement of a set $A \subset \mathcal{X}$. $\bar{A} = \mathcal{X} \setminus A$.
$ A $	The cardinality of a set $A \subset \mathcal{X}$ if it is countable.
\vee, \wedge	Logical <i>Or</i> and <i>And</i> respectively.
$\lesssim, \gtrsim, \asymp$	Inequalities and equality ignoring constant terms.
q_t, r_t	The instantaneous reward and regret respectively. See Equation (2)
$R(\Lambda)$	The cumulative regret after spending capital Λ . See equation 3.
$S(\Lambda)$	The simple regret after spending capital Λ . See second paragraph under equation 3.
$\zeta^{(m)}$	A bound on the maximum difference between $f^{(m)}$ and $f^{(M)}$, $\ f^{(M)} - f^{(m)}\ _\infty \leq \zeta^{(m)}$.
$\mu_t^{(m)}$	The mean of the m^{th} fidelity GP $f^{(m)}$ conditioned on $\mathcal{D}_t^{(m)}$ at time t .
$\kappa_t^{(m)}$	The covariance of the m^{th} fidelity GP $f^{(m)}$ conditioned on $\mathcal{D}_t^{(m)}$ at time t .
$\sigma_t^{(m)}$	The standard deviation of the m^{th} fidelity GP $f^{(m)}$ conditioned on $\mathcal{D}_t^{(m)}$ at time t .
$\mathbf{x}_t, \mathbf{y}_t$	The queried point and observation at time t .
\mathbf{m}_t	The queried fidelity at time t .
$\mathcal{D}_n^{(m)}$	The set of queries at the m^{th} fidelity until time n $\{(\mathbf{x}_t, \mathbf{y}_t)\}_{t:\mathbf{m}_t=m}$.
β_t	The coefficient trading off exploration and exploitation in the UCB. See Theorems 10, 13.
$\varphi_t^{(m)}(x)$	The upper confidence bound provided by the m^{th} fidelity on $f^{(M)}(x)$. $\varphi_t^{(m)}(x) = \mu_{t-1}^{(m)}(x) + \beta_t^{1/2} \sigma_{t-1}^{(m)}(x) + \zeta^{(m)}$.
$\varphi_t(x)$	The combined upper confidence bound provided by all fidelities on $f^{(M)}(x)$. $\varphi_t(x) = \min_m \varphi_t^{(m)}(x)$.
γ	The parameter in MF-GP-UCB for switching between fidelities. See step 2 of Algorithm 1.
\tilde{R}_n	The cumulative regret for the queries after n rounds, $\tilde{R}_n = \sum_{t=1}^n \lambda^{(\mathbf{m}_t)} r_t$. $r_t = f_* + B$ if $\mathbf{m}_t \neq M$ and $r_t = f_* - f^{(M)}(\mathbf{x}_t)$ if $\mathbf{m}_t = M$.
$T_n^{(m)}(\cdot)$	The number of queries at fidelity m at any point in or subset of \mathcal{X} until time n .
$T_n^{(>m)}(\cdot)$	The number of queries at fidelities greater than m at any point in or subset of \mathcal{X} until time n .
n_Λ	Number of plays by a strategy querying only at fidelity M within capital Λ . $n_\Lambda = \lfloor \Lambda / \lambda^{(M)} \rfloor$.
$\Psi_n(A)$	The maximum information gain of a set $A \subset \mathcal{X}$ after n queries in A . See Definition 1.
$\mathcal{X}^{(m)}$	$(\mathcal{X}^{(m)})_{m=1}^M$ is an entirely problem dependent partitioning of \mathcal{X} . See Equation (6).
$\mathcal{H}^{(m)}, \mathcal{H}_\tau^{(m)}$	$(\mathcal{H}^{(m)})_{m=1}^M, (\mathcal{H}_\tau^{(m)})_{m=1}^M$ are partitionings of \mathcal{X} for the discrete and compact cases respectively. See Equations (6) and (12). The analysis of MF-GP-UCB hinges on these partitionings.
$\mathcal{H}_{\tau,n}^{(m)}$	An additional n -dependent inflation of $\mathcal{H}_\tau^{(m)}$. See paragraph under equation (12).
$\hat{\mathcal{H}}^{(m)}, \tilde{\mathcal{H}}^{(m)}$	The arms “above”/“below” $\mathcal{H}^{(m)}$. $\hat{\mathcal{H}}^{(m)} = \bigcup_{\ell=m+1}^M \mathcal{H}^{(\ell)}$, $\tilde{\mathcal{H}}^{(m)} = \bigcup_{\ell=1}^{m-1} \mathcal{H}^{(\ell)}$.
$\hat{\mathcal{H}}_\tau^{(m)}, \tilde{\mathcal{H}}_\tau^{(m)}$	The arms “above”/“below” $\mathcal{H}_\tau^{(m)}$. Defined similar to above.
$\mathcal{X}_g, \mathcal{X}_b$	The good set and bad sets for $M = 2$ fidelity problems. $\mathcal{X}_g = \mathcal{X}^{(2)}$ and $\mathcal{X}_b = \mathcal{X}^{(1)}$.
$\tilde{\mathcal{X}}_{g,\rho}, \tilde{\mathcal{X}}_{b,\rho}$	The inflations of $\mathcal{X}_g, \mathcal{X}_b$ in the discrete case. $\tilde{\mathcal{X}}_{g,\rho} = \{x; f_* - f^{(1)}(x) \leq \zeta^{(1)} + \rho\gamma\} = \mathcal{H}^{(2)}$ and $\tilde{\mathcal{X}}_{b,\rho} = \mathcal{X} \setminus \tilde{\mathcal{X}}_{g,\rho} = \mathcal{H}^{(1)}$.
$\ddot{\mathcal{X}}_{g,\rho,\tau}, \ddot{\mathcal{X}}_{b,\rho,\tau}$	The inflations of $\mathcal{X}_g, \mathcal{X}_b$ in the compact case. $\ddot{\mathcal{X}}_{g,\tau} = \{x; f_* - f^{(1)}(x) \leq \zeta^{(1)} + \max(\tau, \rho\gamma)\} = \mathcal{H}_{\tau,n}^{(2)}$, and $\ddot{\mathcal{X}}_{b,\tau} = \mathcal{X} \setminus \ddot{\mathcal{X}}_{g,\tau} = \mathcal{H}_{\tau,n}^{(1)}$.
$\ddot{\mathcal{X}}_{g,\rho,\tau,n}$	The additional n -dependent inflation of $\ddot{\mathcal{X}}_{g,\rho,\tau}$. $\ddot{\mathcal{X}}_{g,\rho,\tau,n} = \mathcal{H}_{\tau,n}^{(1)}$.
$\Omega_\varepsilon(A)$	The ε -covering number of a subset $A \subset \mathcal{X}$ in the $\ \cdot\ _2$ metric.

A.2. Review of GP-UCB

The following bounds the regret R_n for the GP-UCB algorithm of [Srinivas et al. \(2010\)](#) after n time steps. The algorithm is given in Algorithm 2.

Theorem 4. (Theorems 1 and 2 in [\(Srinivas et al., 2010\)](#)) Let $f \sim \mathcal{GP}(\mathbf{0}, \kappa)$, $f : \mathcal{X} \rightarrow \mathbb{R}$ and κ satisfy Assumption 8. At each query, we have noisy observations $y = f(x) + \epsilon$ where $\epsilon \sim \mathcal{N}(0, \eta^2)$. Denote $C_1 = 8/\log(1 + \eta^{-2})$. Pick $\delta \in (0, 1)$.

- If \mathcal{X} is a finite discrete set, run GP-UCB with $\beta_t = 2 \log \left(\frac{|\mathcal{X}|t^2\pi^2}{6\delta} \right)$. Then,

$$\mathbb{P} \left(\forall n \geq 1, R_n \leq \sqrt{C_1 n \beta_n \Psi_n(\mathcal{X})} \right) \geq 1 - \delta$$

- If $\mathcal{X} = [0, r]^d$, run GP-UCB with $\beta_t = 2 \log \left(\frac{2\pi^2 t^2}{3\delta} \right) + 2d \log \left(t^2 b d r \sqrt{\frac{4ad}{\delta}} \right)$. Then,

$$\mathbb{P} \left(\forall n \geq 1, R_n \leq \sqrt{C_1 n \beta_n \Psi_n(\mathcal{X})} + 2 \right) \geq 1 - \delta$$

Here $\Psi_n(\mathcal{X})$ is the Maximum Information Gain of \mathcal{X} after n queries (see Definition 1).

Algorithm 2 GP-UCB

Input: kernel κ .

For $t = 1, 2, \dots$

- $\mathcal{D}_0 \leftarrow \emptyset, (\mu_0, \sigma_0^2) \leftarrow (\mathbf{0}, \kappa)$.
 - $(\mu_0, \kappa_0) \leftarrow (\mathbf{0}, \kappa)$
 - **for** $t = 1, 2, \dots$
 1. $\mathbf{x}_t \leftarrow \operatorname{argmax}_{x \in \mathcal{X}} \mu_{t-1}(x) + \beta_t^{1/2} \sigma_{t-1}(x)$
 2. $\mathbf{y}_t \leftarrow \text{Query } f \text{ at } \mathbf{x}_t$.
 3. $\mathcal{D}_t = \mathcal{D}_{t-1} \cup \{(\mathbf{x}_t, \mathbf{y}_t)\}$.
 4. Perform Bayesian posterior updates to obtain μ_t, σ_t (See Equation (1)).
-

A.3. Some Ancillary Results

We will use the following results in our analysis. The first is a standard Gaussian concentration result and the second is an expression for the Information Gain in a GP from [Srinivas et al. \(2010\)](#).

Lemma 5 (Gaussian Concentration). Let $Z \sim \mathcal{N}(0, 1)$. Then $\mathbb{P}(Z > \epsilon) \leq \frac{1}{2} \exp(-\epsilon^2/2)$.

Lemma 6 (Mutual Information in GP, [\(Srinivas et al., 2010\)](#) Lemma 5.3). Let $f \sim \mathcal{GP}(\mathbf{0}, \kappa)$, $f : \mathcal{X} \rightarrow \mathbb{R}$ and we observe $y = f(x) + \epsilon$ where $\epsilon \sim \mathcal{N}(0, \eta^2)$. Let A be a finite subset of \mathcal{X} and f_A, y_A be the function values and observations on this set respectively. Using the basic Gaussian properties they show that the mutual information $I(y_A; f_A)$ is,

$$I(y_A; f_A) = \frac{1}{2} \sum_{t=1}^n \log(1 + \eta^{-2} \sigma_{t-1}^2(x_t)).$$

where σ_{t-1}^2 is the posterior variance after observing the first $t-1$ points.

We conclude this section with the following comment on our assumptions in Section 2.

Remark 7 (Validity of the Assumptions A1, A2, A3). It is sufficient to show that when the functions $f^{(m)}$ are sampled from $\mathcal{GP}(\mathbf{0}, \kappa)$, the latter constraints, i.e. $\|f^{(M)}\|_\infty \leq B$ and $\|f^{(M)} - f^{(m)}\|_\infty \leq \zeta^{(m)} \forall m$, occur with positive probability. Then, a generative mechanism would repeatedly sample the $f^{(m)}$'s from the GP and output them when the constraints are satisfied. For finite discrete \mathcal{X} this is straightforward to show. When \mathcal{X} is a compact space, the statement is true for well behaved kernels. For instance, using Assumption 8 (Appendix C) we can establish a high probability bound on the Lipschitz constant of the GP sample $f^{(M)}$. Since for a given $x \in \mathcal{X}$, $\mathbb{P}(-B < f^{(M)}(x) < 0)$ is positive we just need to make sure that the Lipschitz constant is not larger than $B/\text{diam}(\mathcal{X})$. This bounds $\|f^{(M)}\|_\infty < B$. For the latter constraint, since $f^{(M)} - f^{(m)} \sim \mathcal{GP}(\mathbf{0}, 2\kappa)$ is also a GP, the argument follows in an essentially similar fashion.

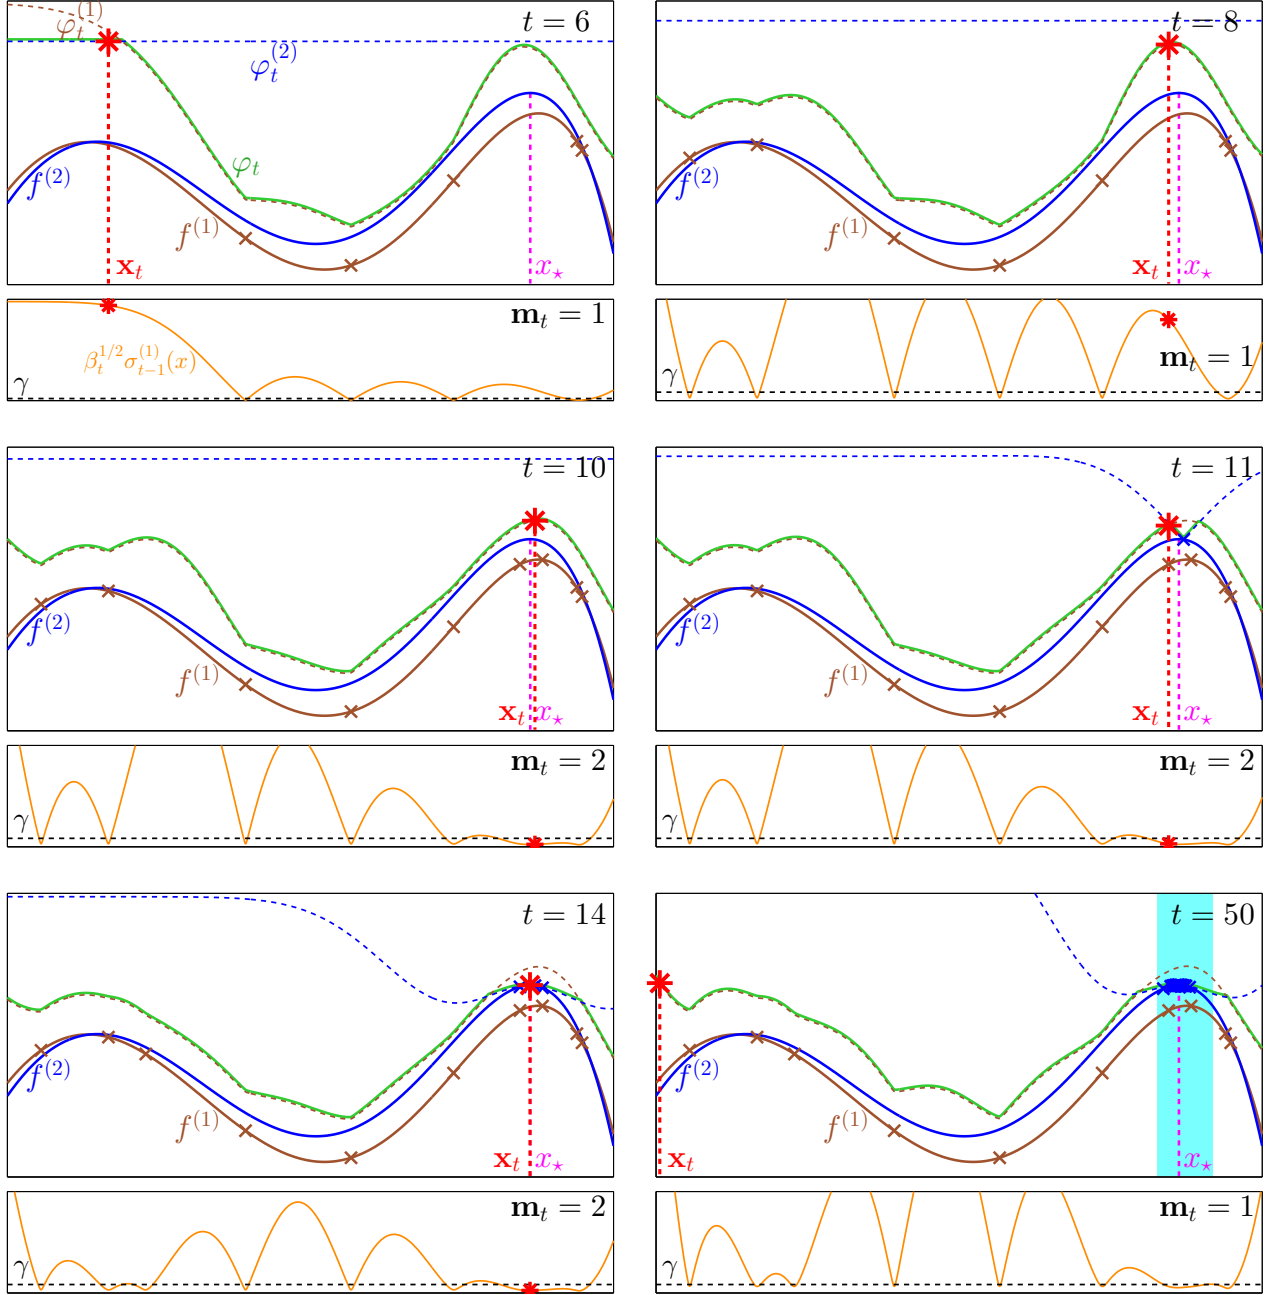


Figure 7. Illustration of MF-GP-UCB for a 2-fidelity problem which has been initialised with 5 random points at the first fidelity. In the top figures, the solid lines in brown and blue are $f^{(1)}, f^{(2)}$ respectively, and the dashed lines are $\varphi_t^{(1)}, \varphi_t^{(2)}$. The solid green line is $\varphi_t = \min(\varphi_t^{(1)}, \varphi_t^{(2)})$. The small crosses are queries from 1 to $t-1$ and the red star is the maximiser of φ_t , i.e. the next query \mathbf{x}_t . x_* , the optimum of $f^{(2)}$ is shown in magenta. In the bottom figures, the solid orange line is $\beta_t^{1/2} \sigma_{t-1}^{(1)}$ and the dashed black line is γ . When $\beta_t^{1/2} \sigma_{t-1}^{(1)}(\mathbf{x}_t) \leq \gamma$ we play at fidelity $\mathbf{m}_t = 2$ and otherwise at $\mathbf{m}_t = 1$.

At the initial stages, MF-GP-UCB is mostly exploring \mathcal{X} in the first fidelity. $\beta_t^{1/2} \sigma_{t-1}^{(1)}$ is large and we are yet to constrain $f^{(1)}$ well to proceed to $m = 2$. At $t = 10$, we have constrained $f^{(1)}$ sufficiently well at a region around the optimum – $\beta_t^{1/2} \sigma_{t-1}^{(1)}(\mathbf{x}_t)$ falls below γ and we query at $\mathbf{m}_t = 2$. Notice that once we do this (at $t = 11$), $\varphi_t^{(2)}$ dips to change φ_t in that region. At $t = 14$, MF-GP-UCB has identified the maximum x_* with just 4 queries to $f^{(2)}$. In the last figure, at $t = 50$, the algorithm decides to explore at a point far away from the optimum. However, this query occurs in the first fidelity since we have not sufficiently constrained $f^{(1)}(\mathbf{x}_t)$ in this region. The key idea is that it is *not necessary* to query such regions at the second fidelity as the first fidelity alone is enough to conclude that it is suboptimal. Herein lies the crux of our method. The region shaded in cyan in the last figure is the good set $\mathcal{X}_g = \{x; f^{(2)}(x_*) - f^{(1)}(x) \leq \zeta^{(1)}\}$ discussed in Section 4. Our analysis predicts that most second fidelity queries in MF-GP-UCB will be confined to this set with high probability and the simulation corroborates this claim. In addition, observe that in a large portion of \mathcal{X} , φ_t is given by $\varphi_t^{(1)}$ except in a small neighborhood around x_* , where it is given by $\varphi_t^{(2)}$.

B. Some Details on MF-GP-UCB

An Extended Simulation

In Figure 7 we provide an extended version of the simulation of Fig. 3 for a 2 fidelity example. Read the caption under the simulation for more details.

More Implementation Details

Data dependent prior: In our experiments, following recommendations in Bull (2011) and Brochu et al. (2010) all GP methods were initialised with uniform random queries using an initialisation capital Λ_0 . For single fidelity methods, we used it at the M^{th} fidelity, whereas for MF-GP-UCB we used $\Lambda_0/2$ at fidelity 1 and $\Lambda_0/2$ at fidelity 2. After initialising the kernel in this manner, we update the kernel every 25 iterations of the method by maximising the GP marginal likelihood.

Choice of β_t : β_t , as specified in Theorems 4, 13 has unknown constants and tends to be too conservative in practice. Following Kandasamy et al. (2015) we use $\beta_t = 0.2d \log(2t)$ which captures the dominant dependencies on d and t .

Initial ζ, γ : We set both ζ, γ to 1% of the range of initial queries and update them as explained in the main text.

Maximising φ_t : To determine \mathbf{x}_t we maximised φ_t using DiRect (Jones et al., 1993). For other GP methods, the EI, PI, GP-UCB acquisition functions were also maximised using DiRect.

MF-GP-UCB was fairly robust to the above choices except when Λ_0 was set too low in which case, all GP methods performed poorly on some experiments.

C. Theoretical Analysis

In this section we present our main theoretical results. While it is self contained, the reader will benefit from first reading the more intuitive discussion in Section 4. The goal in this section is to bound $R(\Lambda)$ for MF-GP-UCB. Recall,

$$R(\Lambda) = \Lambda f_* - \sum_{t=1}^N \lambda^{(m_t)} q_t - \left(\Lambda - \sum_{t=1}^N \lambda^{(m_t)} \right) (-B) = \underbrace{\left(\Lambda - \sum_{t=1}^N \lambda^{(m_t)} \right) (f_* + B)}_{\tilde{r}(\Lambda)} + \underbrace{\sum_{t=1}^N \lambda^{(m_t)} r_t}_{\tilde{R}(\Lambda)},$$

where N is the random number of plays within capital Λ and q_t, r_t are the instantaneous reward and regret as defined in (2). The first term $\tilde{r}(\Lambda)$ is the residual quantity. It is an artefact of the fact that after the $(N+1)^{\text{th}}$ query, the spent capital would have exceeded Λ . It can be bounded by $\tilde{r}(\Lambda) \leq 2B\lambda^{(M)}$ which is typically small. Our analysis will mostly be dealing with the latter term $\tilde{R}(\Lambda)$ for which we will first bound the quantity $\tilde{R}_n = \sum_{t=1}^n \lambda^{(m_t)} r_t$ after n time steps in terms of n . Then, we will bound the random number of plays N within principal Λ . While $N \leq \lfloor \Lambda/\lambda^{(1)} \rfloor$ is a trivial bound, this will be too loose for our purpose. In fact, we will show that after a sufficiently large number of time steps n , with high probability the number of plays at fidelities lower than M will be sublinear in n . Hence $N \in \mathcal{O}(n_\Lambda)$ where $n_\Lambda = \lfloor \Lambda/\lambda^{(M)} \rfloor$ is the number of plays by any algorithm that operates only at the highest fidelity.

Our strategy to bound \tilde{R}_n will be to identify a (possibly disconnected) measurable region of the space \mathcal{Z} which contains x_* and has high value for the payoff function $f^{(M)}(x)$. \mathcal{Z} will be determined by the approximations provided via the lower fidelity evaluations. Denoting $\bar{\mathcal{Z}} = \mathcal{X} \setminus \mathcal{Z}$, we decompose \tilde{R}_n as follows,

$$\tilde{R}_n \leq \underbrace{2B \sum_{m=1}^{M-1} \lambda^{(m)} T_n^{(m)}(\mathcal{X})}_{\tilde{R}_{n,1}} + \underbrace{\lambda^{(M)} \sum_{\substack{t: m_t=M \\ \mathbf{x}_t \in \mathcal{Z}}} \left(f_* - f^{(M)}(\mathbf{x}_t) \right)}_{\tilde{R}_{n,2}} + \underbrace{\lambda^{(M)} \sum_{\substack{t: m_t=M \\ \mathbf{x}_t \in \bar{\mathcal{Z}}}} \left(f_* - f^{(M)}(\mathbf{x}_t) \right)}_{\tilde{R}_{n,3}}. \quad (5)$$

$\tilde{R}_{n,1}$ is the capital spent on the lower fidelity queries for which we receive no reward. $\tilde{R}_{n,2}$ is the regret due to fidelity M queries in \mathcal{Z} and $\tilde{R}_{n,3}$ is due to fidelity M queries outside \mathcal{Z} . To control $\tilde{R}_{n,1}$ we will first bound $T_n^{(m)}(\mathcal{X})$ for $m < M$. This will typically be small containing only $\text{polylog}(n)/\text{poly}(\gamma)$ and $o(n)$ terms. The last two terms can be controlled using the MIGs Ψ_n of $\mathcal{Z}, \bar{\mathcal{Z}}$ respectively (Definition 1). As we will see, $\tilde{R}_{n,2}$ will be the dominant term in n in our final expression since most of the fidelity M queries will be confined to \mathcal{Z} . $T_M^{(n)}(\bar{\mathcal{Z}})$ will be sublinear in n and hence $\tilde{R}_{n,3}$




Figure 8. Illustration of the partition $\mathcal{X}^{(m)}$'s for a $M = 4$ fidelity problem. The sets $\mathcal{J}_0^{(m)}$ are indicated next to their boundaries. $\mathcal{X}^{(1)}, \mathcal{X}^{(2)}, \mathcal{X}^{(3)}, \mathcal{X}^{(4)}$ are shown in green, blue, yellow and red respectively. An illustration for the sets $\mathcal{H}^{(m)}$'s would look similar.

will be of low order. When the lower fidelities allow us to eliminate a large region of the space, $\text{vol}(\mathcal{Z}) \ll \text{vol}(\overline{\mathcal{Z}})$ and consequently the maximum information gain of \mathcal{Z} will be much smaller than that of $\overline{\mathcal{Z}}$, $\Psi_n(\mathcal{Z}) \ll \Psi_n(\overline{\mathcal{Z}})$. As we will see, this results in much better regret for MF-GP-UCB in comparison to GP-UCB.

For the analysis, we will need the following regularity conditions on the kernel. It is satisfied for four times differentiable kernels such as the SE and Matérn kernels with smoothness parameter $\nu > 2$ (Ghosal & Roy, 2006).

Assumption 8. Let $f \sim \mathcal{GP}(\mathbf{0}, \kappa)$, where $\kappa : [0, r]^d \times [0, r]^d \rightarrow \mathbb{R}$ is a stationary kernel. The partial derivatives of f satisfies the following high probability bound. There exists constants $a, b > 0$ such that, for all $J > 0$,

$$\forall i \in \{1, \dots, d\}, \quad \mathbb{P} \left(\sup_x \left| \frac{\partial f(x)}{\partial x_i} \right| > J \right) \leq ae^{-(J/b)^2}.$$

We will need this assumption even on discrete \mathcal{X} – one interpretation is that the GP was sampled on $[0, r]^d$ but we are only operating on a discrete subset. For the proofs of both the discrete and compact cases, we will need to control the conditional variances for queries within a subset $A \subset \mathcal{X}$. To that end, we provide the lemma below.

Lemma 9. Let $f \sim \mathcal{GP}(\mathbf{0}, \kappa)$, $f : \mathcal{X} \rightarrow \mathbb{R}$ and each time we query at any $x \in \mathcal{X}$ we observe $y = f(x) + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \eta^2)$. Let $A \subset \mathcal{X}$. Assume that we have queried f at n points, $(x_t)_{t=1}^n$ of which s points are in A . Let σ_{t-1} denote the posterior variance at time t , i.e. after $t - 1$ queries. Then, $\sum_{x_t \in A} \sigma_{t-1}^2(x_t) \leq \frac{2}{\log(1+\eta^{-2})} \Psi_s(A)$.

Proof Let $A_s = \{z_1, z_2, \dots, z_s\}$ be the queries inside A in the order they were queried. Now, assuming that we have only queried inside A at A_s , denote by $\tilde{\sigma}_{t-1}(\cdot)$, the posterior standard deviation after $t - 1$ such queries. Then,

$$\sum_{t: x_t \in A} \sigma_{t-1}^2(x_t) \leq \sum_{t=1}^s \tilde{\sigma}_{t-1}^2(z_t) \leq \sum_{t=1}^s \eta^2 \frac{\tilde{\sigma}_{t-1}^2(z_t)}{\eta^2} \leq \sum_{t=1}^s \frac{\log(1 + \eta^{-2} \tilde{\sigma}_{t-1}^2(z_t))}{\log(1 + \eta^{-2})} \leq \frac{2}{\log(1 + \eta^{-2})} I(y_{A_s}; f_{A_s})$$

Queries outside A will only decrease the variance of the GP so we can upper bound the first sum by the posterior variances of the GP with only the queries in A . The third step uses the inequality $u^2/v^2 \leq \log(1 + u^2)/\log(1 + v^2)$ with $u = \tilde{\sigma}_{t-1}(z_t)/\eta$ and $v = 1/\eta$ and the last step uses Lemma 6. The result follows from the fact that $\Psi_s(A)$ maximises the mutual information among all subsets of size s . ■

C.1. Discrete \mathcal{X}

We first analyse when \mathcal{X} is a discrete subset of $[0, r]^d$. Denote $\Delta^{(m)}(x) = f_\star - f^{(m)}(x) - \zeta^{(m)}$ and $\mathcal{J}_\eta^{(m)} = \{x \in \mathcal{X}; \Delta^{(m)}(x) \leq \eta\}$. Let $\rho > 1$ be given. Central to our analysis will be two partitionings $(\mathcal{X}^{(m)})_{m=1}^M, (\mathcal{H}^{(m)})_{m=1}^M$ of \mathcal{X} . The latter depends on the parameter γ and the given ρ . Let $\mathcal{X}^{(1)} = \overline{\mathcal{J}}_0^{(1)}, \mathcal{H}^{(1)} = \overline{\mathcal{J}}_{\rho\gamma}^{(1)}$. Then define,

$$\begin{aligned} \mathcal{X}^{(m)} &= \overline{\mathcal{J}}_0^{(m)} \cap \left(\bigcap_{\ell=1}^{m-1} \mathcal{J}_0^{(\ell)} \right) \quad \text{for } 2 \leq m \leq M-1, & \mathcal{X}^{(M)} &= \bigcap_{\ell=1}^{M-1} \mathcal{J}_0^{(\ell)}. \\ \mathcal{H}^{(m)} &= \overline{\mathcal{J}}_{\rho\gamma}^{(m)} \cap \left(\bigcap_{\ell=1}^{m-1} \mathcal{J}_{\rho\gamma}^{(\ell)} \right) \quad \text{for } 2 \leq m \leq M-1, & \mathcal{H}^{(M)} &= \bigcap_{\ell=1}^{M-1} \mathcal{J}_{\rho\gamma}^{(\ell)}. \end{aligned} \quad (6)$$

We have illustrated these sets in figure 8. The partitioning $\{\mathcal{X}^{(m)}\}_{m=1}^M$ is entirely problem dependent and is fundamental to the problem. Intuitively, $\mathcal{X}^{(m)}$ is the set of points in \mathcal{X} that can be safely excluded from queries beyond the m^{th} fidelity

since they are sufficiently far away from f_* . We are interested in settings where $|\mathcal{X}^{(m)}| \gg |\mathcal{X}^{(m-1)}|$ but $\lambda^{(m)} \ll \lambda^{(m-1)}$. For instance consider an $M = 2$ problem in optimal policy search in robotics, where the first fidelity is a simulator and the second fidelity is the real world experiment. Hence $\lambda^{(1)} \ll \lambda^{(2)}$. The simulator can be used to exclude several bad policies and we can use the expensive real world experiments on a promising set of candidates $\mathcal{X}^{(2)}$ as informed by the simulations.

Our analysis of MF-GP-UCB will use slightly inflated regions $\mathcal{H}^{(m)}$ instead of $\mathcal{X}^{(m)}$. The parameter γ causes us to explore at a slightly larger region than $\mathcal{X}^{(m)}$ at the m^{th} fidelity. $\mathcal{X}^{(m)}$ is the smallest such set that can be excluded from $> m$ fidelity queries by letting $\gamma \rightarrow 0$. In addition to the above, we will also find it useful to define the sets “above” $\mathcal{H}^{(m)}$ as $\widehat{\mathcal{H}}^{(m)} = \bigcup_{\ell=m+1}^M \mathcal{H}^{(\ell)}$ and the sets “below” $\mathcal{H}^{(m)}$ as $\widetilde{\mathcal{H}}^{(m)} = \bigcup_{\ell=1}^{m-1} \mathcal{H}^{(\ell)}$. Intuitively, $\mathcal{H}^{(m)}$ is the set of points that MF-GP-UCB will query at the m^{th} fidelity but exclude from higher fidelities due to information from fidelity m . $\widetilde{\mathcal{H}}^{(m)}$ is the set of points that can be excluded from queries at fidelities m and beyond due to information from lower fidelities. $\widehat{\mathcal{H}}^{(m)}$ are points that need to be queried at fidelities higher than m . In the 2 fidelity setting described in Section 4, the bad set \mathcal{X}_b is $\mathcal{X}^{(1)}$ and the good set \mathcal{X}_g is $\mathcal{X}^{(2)}$. Similarly, $\mathcal{X}_{b,\rho}$ is $\mathcal{H}^{(1)}$ and $\mathcal{X}_{g,\rho}$ is $\mathcal{H}^{(2)}$.

We now state our main theorem for the discrete case.

Theorem 10. *Let \mathcal{X} be a discrete subset of $[0, r]^d$. Let $f^{(m)} \sim \mathcal{GP}(\mathbf{0}, \kappa)$ for all m . Assume that $f^{(m)}$ ’s satisfy assumptions A2, A3 and κ satisfy Assumption 8. Pick $\delta \in (0, 1)$ and run MF-GP-UCB with $\beta_t = 2 \log \left(\frac{M|\mathcal{X}|\pi^2 t^2}{3\delta} \right)$. Then, for all $\alpha \in (0, 1)$, and $\rho = \max(2, 1 + \sqrt{1/2 + 1/\alpha})$ and sufficiently large Λ we have $R(\Lambda) \in \mathcal{O} \left(\lambda^{(M)} \sqrt{n_\Lambda \beta_{n_\Lambda} \Psi_{n_\Lambda}(\mathcal{H}^{(M)})} + \lambda^{(M)} \sqrt{n_\Lambda^\alpha \beta_{n_\Lambda} \Psi_{n_\Lambda^\alpha}(\mathcal{X})} \right)$, with high probability. Here $n_\Lambda = \lfloor \Lambda / \lambda^{(M)} \rfloor$ is the number of plays by an algorithm that only plays at the highest fidelity. Precisely, there exists Λ_0 such that for all $\Lambda \geq \Lambda_0$ with probability $> 1 - \delta$ we have,*

$$R(\Lambda) \leq 2B\lambda^{(M)} + 2B \sum_{m=1}^{M-1} \lambda^{(m)} \left(\sum_{x \in \mathcal{H}^{(m)}} \left\{ \frac{5\eta^2}{\Delta^{(m)}(x)^2} \beta_{2n_\Lambda} + 1 \right\} + |\widetilde{\mathcal{H}}^{(m)}| (2n_\Lambda)^\alpha + |\widehat{\mathcal{H}}^{(m)}| \left(\frac{\eta^2}{\gamma^2} \beta_{2n_\Lambda} + 1 \right) \right) \\ + \lambda^{(M)} \sqrt{2C_1 n_\Lambda \beta_{2n_\Lambda} \Psi_{2n_\Lambda}(\mathcal{H}^{(M)})} + \lambda^{(M)} \sqrt{2|\mathcal{X}| C_1 n_\Lambda^\alpha \beta_{2n_\Lambda} \Psi_{2|\mathcal{X}| n_\Lambda^\alpha}(\mathcal{X})},$$

where $C_1 = 8 / \log(1 + \eta^2)$. Recall that the sets $\mathcal{H}^{(m)}$ depend on ρ .

Synopsis: Before we provide the proof, let us interpret the terms above in order to distill the gist of the theorem. From Theorem 4 we can infer that the regret for GP-UCB which plays only at the highest fidelity is $\mathcal{O} \left(\lambda^{(M)} \sqrt{n_\Lambda \beta_{n_\Lambda} \Psi_{n_\Lambda}(\mathcal{X})} \right)$. Ignoring the common terms and constants, the regret for MF-GP-UCB is $\lambda^{(M)} \sqrt{n_\Lambda \beta_{n_\Lambda} \Psi_{n_\Lambda}(\mathcal{H}^{(M)})} + \lambda^{(M)} \sqrt{n_\Lambda^\alpha \beta_{n_\Lambda} \Psi_{n_\Lambda^\alpha}(\mathcal{X})}$, $\forall \alpha > 0$ and that for GP-UCB is $\lambda^{(M)} \sqrt{n_\Lambda \beta_{n_\Lambda} \Psi_{n_\Lambda}(\mathcal{X})}$ for GP-UCB. In the cases where $\text{vol}(\mathcal{H}^{(M)}) \ll \text{vol}(\mathcal{X})$, i.e. the lower fidelities provide a good approximation for $f^{(M)}$, MF-GP-UCB significantly outperforms GP-UCB whose exploration complexity Ψ for \mathcal{X} is coupled with n_Λ . In contrast, the coupling with n_Λ for MF-GP-UCB is with the much smaller region $\mathcal{H}^{(M)}$ and the MIG Ψ_n for \mathcal{X} is only sublinear in n_Λ . This directly captures our intuitions behind using the lower fidelities: to eliminate the bad regions of \mathcal{X} and use the high fidelity evaluations only to explore at the promising regions. The price we have to pay though is that we spend part of our capital on lower fidelity evaluations. But these terms are either logarithmic or sublinear in n_Λ . In addition, these evaluations are cheaper than fidelity M evaluations so these terms are of low order.

Proof of Theorem 10. First we consider MF-GP-UCB after n steps regardless of which fidelities it has queried at and obtain a bound on \tilde{R}_n . Then we will bound the number of plays N within capital Λ . We will need the following two lemmas for our analysis. Lemma 11 is used to establish that $\varphi_t(x)$ upper bounds $f^{(M)}(x)$. Lemma 12 bounds $T_n^{(m)}(x)$ for different x, m which we will use to control $\tilde{R}_{n,1}$ and $\tilde{R}_{n,3}$. The proofs are given in Sections C.1.1 and C.1.2.

Lemma 11. *Pick $\delta \in (0, 1)$ and choose $\beta_t = 2 \log \left(\frac{M|\mathcal{X}|\pi^2 t^2}{3\delta} \right)$. Then, with probability at least $1 - \delta/2$, for all $t \geq 1$, for all $x \in \mathcal{X}$ and for all $m \in \{1, \dots, M\}$, we have $|f^{(m)}(x) - \mu_{t-1}^{(m)}(x)| \leq \beta_t^{1/2} \sigma_{t-1}^{(m)}(x)$.*

Lemma 12. *Pick $\delta \in (0, 1)$ and set $\beta_t = 2 \log \left(\frac{M|\mathcal{X}|\pi^2 t^2}{3\delta} \right)$. Let $\rho \geq 2$. Further assume $\varphi_t(x_*) \geq f_*$. Consider any*

	$\mathcal{H}^{(1)}$	$\mathcal{H}^{(2)}$	$\mathcal{H}^{(m)}$		$\mathcal{H}^{(M)} \setminus \{x_\star\}$		
$T_n^{(1)}(x)$	n^α	$\frac{\eta^2}{\gamma^2} \beta_n + 1$	\dots	$\frac{\eta^2}{\gamma^2} \beta_n + 1$	\dots	$\frac{\eta^2}{\gamma^2} \beta_n + 1$	
$T_n^{(2)}(x)$		$\frac{5\eta^2}{\Delta^{(m)}(x)^2} \beta_n + 1$	\dots	$\frac{\eta^2}{\gamma^2} \beta_n + 1$	\dots	$\frac{\eta^2}{\gamma^2} \beta_n + 1$	
\vdots		n^α	\dots		$\frac{5\eta^2}{\Delta^{(m)}(x)^2} \beta_n + 1$	\dots	$\frac{\eta^2}{\gamma^2} \beta_n + 1$
$T_n^{(m)}(x)$							
\vdots							
$T_n^{(M)}(x)$	n^α				$\frac{5\eta^2}{\Delta^{(m)}(x)^2} \beta_n + 1$		

Table 1. Bounds on the number of queries for each $x \in \mathcal{H}^{(m)}$ (columns) at each fidelity (rows). The bound for $T_n^{(M)}(x)$ in $\mathcal{H}^{(M)}$ holds for all arms except the optimal arm x_\star (note $\Delta^{(M)}(x_\star) = 0$).

$x \in \mathcal{H}^{(m)} \setminus \{x_\star\}$ for $m < M$. We then have the following bounds on the number of queries at any given time step n ,

$$\begin{aligned}
 T_n^{(\ell)}(x) &\leq \frac{\eta^2}{\gamma^2} \beta_n + 1, \quad \text{for } \ell < m, \\
 \mathbb{P} \left(T_n^{(m)}(x) > \left\lceil 5 \left(\frac{\eta}{\Delta^{(m)}(x)} \right)^2 \beta_n \right\rceil \right) &\leq \frac{3\delta}{2\pi^2} \frac{1}{|\mathcal{X}|n^2}, \\
 \mathbb{P} \left(T_n^{(>m)}(x) > u \right) &\leq \frac{3\delta}{2\pi^2} \frac{1}{|\mathcal{X}|u^{2(\rho-1)^2-1}}.
 \end{aligned}$$

First whenever $\varphi_t(x_\star) \geq f_\star$, by using the union bound on the second result of Lemma 12, we have

$$\mathbb{P} \left(\forall n \geq 1, \forall m \in \{1, \dots, M\}, \forall x \in \mathcal{H}^{(m)} \setminus \{x_\star\}, T_n^{(m)}(x) > \left\lceil 5 \left(\frac{\eta}{\Delta^{(m)}(x)} \right)^2 \beta_n \right\rceil \right) \leq \frac{\delta}{4}.$$

Here we have used $\sum n^{-2} = \pi^2/6$ and that the last two quantifiers just enumerates over all $x \in \mathcal{X} \setminus \{x_\star\}$. In the third result of Lemma 12 we use $u = n^\alpha$ and $\rho = 1 + \sqrt{1/2 + 1/\alpha}$ as given in the theorem. Then, $u^{2(\rho-1)^2-1} = n^2$. Applying the union bound once again we have,

$$\mathbb{P} \left(\forall n \geq 1, \forall m \in \{1, \dots, M\}, \forall x \in \mathcal{H}^{(m)}, T_n^{(>m)}(x) > n^\alpha \right) \leq \frac{\delta}{4}$$

The condition for Lemma 12 holds with probability $> 1 - \delta/2$ (by Lemma 11), and the above bounds hold together with probability $> 1 - \delta$. We have tabulated the bounds in Table 1 and will use it to bound \tilde{R}_n . Consider $T_n^{(m)}(x)$ for any $m < M$. From the table we can read off the following bound,

$$T_n^{(m)}(\mathcal{X}) \leq |\tilde{\mathcal{H}}^{(m)}| n^\alpha + |\hat{\mathcal{H}}^{(m)}| \left(\frac{\eta^2}{\gamma^2} \log(n) + 1 \right) + \sum_{x \in \mathcal{H}^{(m)}} \frac{5\eta^2}{\Delta^{(m)}(x)^2} \beta_n + 1$$

By summing the above over $1, \dots, M-1$ we can bound $\tilde{R}_{n,1}$. Now we use Lemma 11 to control the instantaneous regret.

$$\begin{aligned}
 f_\star - f^{(M)}(\mathbf{x}_t) &\leq \varphi_t(x_\star) - (\mu_{t-1}^{(M)}(\mathbf{x}_t) - \beta_t^{1/2} \sigma_{t-1}^{(M)}(\mathbf{x}_t)) \leq \varphi_t(\mathbf{x}_t) - (\mu_{t-1}^{(M)}(\mathbf{x}_t) - \beta_t^{1/2} \sigma_{t-1}^{(M)}(\mathbf{x}_t)) \\
 &\leq 2\beta_t^{1/2} \sigma_{t-1}^{(M)}(\mathbf{x}_t)
 \end{aligned} \tag{7}$$

$\varphi_t^{(m)}(x)$ is an upper bound for $f^{(M)}(x)$ by Lemma 11 and the assumption A2, and hence so is the minimum $\varphi_t(x)$. Above we have used that \mathbf{x}_t was the maximiser of $\varphi_t(x)$ and that $\varphi_t^{(M)}(x) \geq \varphi_t(x)$. To control $\tilde{R}_{n,2}, \tilde{R}_{n,3}$ we will use $\mathcal{Z} = \mathcal{H}^{(M)}$ and $\bar{\mathcal{Z}} = \bar{\mathcal{H}}^{(M)} = \tilde{\mathcal{H}}^{(M)}$ and invoke Lemma 9. Applying Jensen's inequality in the form $(\sum_{i=1}^s a_i)^2 \leq s \sum_{i=1}^s a_i^2$ yields,

$$\begin{aligned}
 \tilde{R}_{n,3}^2 &\leq T_n^{(m)}(\bar{\mathcal{Z}}) \sum_{\substack{t: \mathbf{m}_t = M \\ \mathbf{x}_t \in \bar{\mathcal{Z}}}} \left(f_\star - f^{(M)}(\mathbf{x}_t) \right)^2 \leq T_n^{(m)}(\bar{\mathcal{Z}}) \sum_{\substack{t: \mathbf{m}_t = M \\ \mathbf{x}_t \in \bar{\mathcal{Z}}}} 4\beta_t(\sigma_{t-1}^{(m)}(\mathbf{x}_t))^2 \\
 &\leq C_1 T_n^{(m)}(\bar{\mathcal{Z}}) \beta_n \Psi_{T_n^{(m)}(\bar{\mathcal{Z}})}(\bar{\mathcal{Z}})
 \end{aligned} \tag{8}$$

where $C_1 = 8/\log(1+\eta^{-2})$. Once again, using Table 1 we can bound $T_n^{(M)}(\tilde{\mathcal{H}}^{(M)}) \leq |\tilde{\mathcal{H}}^{(M)}|n^\alpha \leq |\mathcal{X}|n^\alpha$. Therefore we have $\tilde{R}_{n,3} \leq \sqrt{|\mathcal{X}|^\alpha C_1 n^\alpha \beta_n \Psi_{(|\mathcal{X}|n)^\alpha}(\mathcal{X})}$. Bounding $\tilde{R}_{n,2}$ follows essentially a similar procedure. Since $T_n^{(m)}(\mathcal{H}^{(m)}) \leq n$ trivially, we have $\tilde{R}_{n,2} \leq \sqrt{C_1 n \beta_n \Psi_n(\mathcal{X})}$.

The expressions we now have for $\tilde{R}_{n,1}, \tilde{R}_{n,2}, \tilde{R}_{n,3}$ match the expressions given in the theorem by replacing $2n_\Lambda$ with n . Next we will show that for sufficiently large Λ , $N \leq 2n_\Lambda$ which will complete the proof. For this we first observe,

$$\sum_{m=1}^{M-1} \sum_{x \in \mathcal{H}^{(m)}} T_n^{(m)}(x) = \sum_{m=1}^{M-1} \sum_{\ell=1}^{m-1} \sum_{x \in \mathcal{H}^{(\ell)}} T_n^{(m)}(x) \leq \sum_{m=1}^{M-2} \sum_{x \in \mathcal{H}^{(m)}} T_n^{(>m)}(x) \leq |\mathcal{X}|n^\alpha. \quad (9)$$

Therefore, the following upper bounds the number of queries at fidelities less than M after n time steps,

$$|\mathcal{X}|n^\alpha + \sum_{m=1}^{M-1} \sum_{x \in \mathcal{H}^{(m)}} \left(\frac{5\eta^2}{\Delta^{(m)}(x)^2} \beta_n + 1 \right) + \sum_{m=1}^{M-1} |\hat{\mathcal{H}}^{(m)}| \left(\frac{\eta^2}{\gamma^2} \beta_n + 1 \right).$$

Note that since $\beta_n \in \mathcal{O}(\log(n))$ this bound is sublinear in n . Say we have queried an n_0 number of times such that for all $n \geq n_0$, the above is less than $n/2$. For all such n , $T_n^{(M)}(\mathcal{X}) > n/2$ and therefore the expended budget after n rounds $\Lambda(n)$ satisfies $\Lambda(n) \geq \lambda^{(M)}n/2$. Since our bounds hold with probability $> 1 - \delta$ for all n we can invert the above inequality to bound the number of random plays N after some capital Λ ; $N \leq 2\Lambda/\lambda^{(M)} = 2n_\Lambda$. We only need to make sure that $N \geq n_0$ which can be guaranteed if $\Lambda \geq \Lambda_0 := \lambda^{(M)}(n_0 + 1)$ since $N \geq \lfloor \Lambda/\lambda^{(M)} \rfloor$. Therefore $R(\Lambda) \leq \tilde{r}(\Lambda) + \tilde{R}_{2n_\Lambda,1} + \tilde{R}_{2n_\Lambda,2} + \tilde{R}_{2n_\Lambda,3}$. The theorem follows with the final detail $\tilde{r}(\Lambda) \leq 2B\lambda^{(M)}$. ■

C.1.1. PROOF OF LEMMA 11

This is a straightforward argument using Gaussian concentration and the union bound but we present it here to point out a subtle conditioning argument with multiple fidelities. First consider any given m, t, x .

$$\begin{aligned} \mathbb{P}\left(|f^{(m)}(x) - \mu_{t-1}^{(m)}(x)| > \beta_t^{1/2} \sigma_{t-1}^{(m)}(x)\right) &= \mathbb{E}\left[\mathbb{E}\left[|f^{(m)}(x) - \mu_{t-1}^{(m)}(x)| > \beta_t^{1/2} \sigma_{t-1}^{(m)}(x) \mid \mathcal{D}_{t-1}^{(m)}\right]\right] \\ &= \mathbb{E}\left[\mathbb{P}_{Z \sim \mathcal{N}(0,1)}\left(|Z| > \beta_t^{1/2}\right)\right] \leq \exp\left(-\frac{\beta_t}{2}\right) = \frac{3\delta}{M|\mathcal{X}|\pi^2 t^2}. \end{aligned}$$

In the first step we have conditioned w.r.t $\mathcal{D}_{t-1}^{(m)}$ which allows us to use Lemma 5. The posterior conditioned on all queries will not be a Gaussian due to the $\zeta^{(m)}$ constraints. We will be using this conditioning argument repeatedly in our analysis. The statement follows via a union bound over all $m \in \{1, \dots, M\}$, $x \in \mathcal{X}$ and all t and noting that $\sum_t t^{-2} = \pi^2/6$. ■

C.1.2. PROOF OF LEMMA 12

First consider any $\ell < m$. Assume that we have already queried $\eta^2 \beta_n / \gamma^2 + 1$ times at any $t \leq n$. Since the Gaussian variance after s observations is η^2/s and that queries elsewhere will only decrease the conditional variance we have, $\kappa_{t-1}^{(\ell)}(x, x) \leq \frac{\eta^2}{T_{t-1}^{(\ell)}(x)} < \frac{\gamma^2}{\beta_n}$. Therefore, $\beta_t^{1/2} \sigma_{t-1}^{(\ell)}(x) < \beta_n^{1/2} \sigma_{t-1}^{(\ell)}(x) < \gamma$ and by the design of our algorithm we will not play at the ℓ^{th} fidelity at time t for all t until n . This establishes the first result.

To bound $T_n^{(m)}(x)$ we first observe,

$$\begin{aligned} \mathbb{1}\{T_n^{(m)}(x) > u\} &\leq \mathbb{1}\{\exists t : u+1 \leq t \leq n : \varphi_t(x) \text{ was maximum} \quad \wedge \quad \beta_t^{1/2} \sigma_{t-1}^{(\ell)}(x) < \gamma, \forall \ell < m \quad \wedge \\ &\quad \beta_t^{1/2} \sigma_{t-1}^{(m)}(x) \geq \gamma \quad \wedge \quad T_{t-1}^{(m)}(x) \geq u\} \\ &\leq \mathbb{1}\{\exists t : u+1 \leq t \leq n : \varphi_t(x) > \varphi_t(x_\star) \quad \wedge \quad T_{t-1}^{(m)}(x) \geq u\} \\ &\leq \mathbb{1}\{\exists t : u+1 \leq t \leq n : \varphi_t^{(m)}(x) > f_\star \quad \wedge \quad T_{t-1}^{(m)}(x) \geq u\}. \end{aligned} \quad (10)$$

The first line just enumerates the conditions in our algorithm for it to have played x at time t at fidelity m . In the second step we have relaxed some of those conditions, noting in particular that if $\varphi_t(\cdot)$ was maximised at x then it must be larger than $\varphi_t(x_\star)$. The last step uses the fact that $\varphi_t^{(m)}(x) \geq \varphi_t(x)$ and the assumption on $\varphi_t(x_\star)$. Consider the event

$\{\varphi_t^{(m)}(x) > f_\star \wedge T_{t-1}^{(m)}(x) \geq u\}$. We will choose $u = \lceil 5\eta^2 \beta_n / \Delta^{(m)}(x)^2 \rceil$ and bound its probability via,

$$\begin{aligned} \mathbb{P}(\varphi_t^{(m)}(x) > f_\star \wedge T_{t-1}^{(m)}(x) \geq u) &= \mathbb{P}(\mu_{t-1}^{(m)}(x) + \beta_t^{1/2} \sigma_{t-1}^{(m)}(x) + \zeta^{(m)} > f_\star \wedge T_{t-1}^{(m)}(x) \geq u) \\ &= \mathbb{P}\left(\mu_{t-1}^{(m)}(x) - f^{(m)}(x) > \underbrace{f_\star - f^{(m)}(x) - \zeta^{(m)}}_{\Delta^{(m)}(x)} - \beta_t^{1/2} \sigma_{t-1}^{(m)}(x) \wedge T_{t-1}^{(m)}(x) \geq u\right) \\ &\leq \mathbb{P}\left(\mu_{t-1}^{(m)}(x) - f^{(m)}(x) > (\sqrt{5} - 1)\beta_n^{1/2} \sigma_{t-1}^{(m)}(x)\right) \\ &\leq \mathbb{P}_{Z \sim \mathcal{N}(0,1)}\left(Z > \frac{(\sqrt{5} - 1)^2}{2} \beta_n^{1/2}\right) \leq \frac{1}{2} \exp\left(-\frac{3}{4} \beta_n\right) = \frac{1}{2} \left(\frac{3\delta}{M|\mathcal{X}|\pi^2}\right)^{\frac{3}{2}} n^{-3} \leq \frac{1}{2} \frac{3\delta}{|\mathcal{X}|\pi^2} n^{-3} \end{aligned}$$

Above in the third step we have used, if $u \geq 5\eta^2 \beta_n / \Delta^{(m)}(x)^2$, then $\Delta^{(m)}(x) \geq \sqrt{5} \beta_n^{1/2} \sigma_{t-1}^{(m)}(x)$ and that $\beta_n \geq \beta_t$. The fourth step uses Lemma 5 after conditioning on $\mathcal{D}_{t-1}^{(m)}$, the fifth step uses $(\sqrt{5} - 1)^2 > 3/2$ and the last step uses $3\delta/|\mathcal{X}|\pi^2 < 1$. Now use the union bound on (10), $\mathbb{P}(T_n^{(m)}(x) > u) \leq \sum_{t=u+1}^n \mathbb{P}(\varphi_t^{(m)}(x) > f_\star \wedge T_{t-1}^{(m)}(x) \geq u)$. The second inequality of the theorem follows by noting that there are at most n terms in the summation.

Finally, for the third inequality we observe

$$\mathbb{P}(T_n^{(>m)}(x) > u) \leq \mathbb{P}(\exists t : u+1 \leq t \leq n; \varphi_t^{(m)}(x) > f_\star \wedge \beta_t^{1/2} \sigma_{t-1}^{(m)}(x) < \gamma). \quad (11)$$

As before, we have used that if x is to be queried at time t , then $\varphi_t(x)$ should be at least larger than $\varphi_t(x_\star)$ which is larger than f_\star due to the assumption in the theorem. The second condition is necessary to ensure that the switching procedure proceeds beyond the m^{th} fidelity. It is also necessary to have $\beta_t^{1/2} \sigma_{t-1}^{(\ell)}(x) < \gamma$ for $\ell < m$, but we have relaxed them. We first bound the probability of the event $\{\varphi_t^{(m)}(x) > f_\star \wedge \beta_t^{1/2} \sigma_{t-1}^{(m)}(x) < \gamma\}$.

$$\begin{aligned} \mathbb{P}(\varphi_t^{(m)}(x) > f_\star \wedge \beta_t^{1/2} \sigma_{t-1}^{(m)}(x) < \gamma) &= \mathbb{P}(\mu_{t-1}^{(m)}(x) - f^{(m)}(x) > \Delta^{(m)}(x) - \beta_t^{1/2} \sigma_{t-1}^{(m)}(x) \wedge \beta_t^{1/2} \sigma_{t-1}^{(m)}(x) < \gamma) \\ &\leq \mathbb{P}(\mu_{t-1}^{(m)}(x) - f^{(m)}(x) > \rho\gamma - \beta_t^{1/2} \sigma_{t-1}^{(m)}(x) \wedge \beta_t^{1/2} \sigma_{t-1}^{(m)}(x) < \gamma) \\ &\leq \mathbb{P}(\mu_{t-1}^{(m)}(x) - f^{(m)}(x) > (\rho - 1)\beta_t^{1/2} \sigma_{t-1}^{(m)}(x)) \leq \mathbb{P}_{Z \sim \mathcal{N}(0,1)}\left(Z > (\rho - 1)\beta_t^{1/2}\right) \\ &\leq \frac{1}{2} \exp\left(-\frac{(\rho - 1)^2}{2} \beta_t\right) = \frac{1}{2} \left(\frac{3\delta}{M|\mathcal{X}|\pi^2}\right)^{(\rho-1)^2} t^{-2(\rho-1)^2} \leq \frac{1}{2} \frac{3\delta}{|\mathcal{X}|\pi^2} t^{-2(\rho-1)^2} \end{aligned}$$

Here, the second step uses that for all $x \in \mathcal{H}^{(m)}$, $\Delta^{(m)}(x) > \rho\gamma$. The third step uses the second condition and the last step uses that $\rho \geq 2$. Using the union bound on (11) and bounding the sum by an integral gives us,

$$\mathbb{P}(T_n^{(>m)}(x) > u) \leq \sum_{t=u+1}^n \frac{1}{2} \frac{3\delta}{|\mathcal{X}|\pi^2} t^{-2(\rho-1)^2} \leq \frac{1}{2} \frac{3\delta}{|\mathcal{X}|\pi^2} \int_u^\infty t^{-2(\rho-1)^2} dt \leq \frac{1}{2} \frac{3\delta}{|\mathcal{X}|\pi^2} \frac{1}{u^{2(\rho-1)^2-1}}. \quad \blacksquare$$

Some Remarks

1. Since the switching criterion of MF-GP-UCB (step 2 in Algorithm 1) makes sure that we do not query often at a lower fidelity, an easy upper bound for $T_n^{(m)}(x)$ is $\frac{\eta^2}{\gamma^2} \beta_n + 1$. While we have used this to bound the number of plays in $\widehat{\mathcal{H}}^{(m)}$ we have used different techniques for the points in $\mathcal{H}^{(m)}, \widetilde{\mathcal{H}}^{(m)}$. The main reason is to avoid the γ^{-2} dependence on these arms, especially since we are interested in settings where $|\widetilde{\mathcal{H}}^{(m)}| \gg |\mathcal{H}^{(m)}| \gg |\widehat{\mathcal{H}}^{(m)}|$. In particular, we can control the number of plays at the m^{th} fidelity for any $x \in \mathcal{H}^{(m)}$ using the difference $\Delta^{(m)}(x)$. For arms that are far away from the optimum this provides a tighter bound. Even though the dependence is polynomial in n for $\widetilde{\mathcal{H}}^{(m)}$, its best to avoid the γ^{-2} dependence since we would like to keep γ as small as possible so that $\mathcal{H}^{(m)}$ approaches $\mathcal{X}^{(m)}$.
2. A polylog(n_Λ) rate can also be shown for the discrete case, similar to Theorem 1 in Dani et al. (2008), but it will have worse dependence on Ψ . We stick to the above form which is comparable to the results of Srinivas et al. (2010).

C.2. Compact and Convex \mathcal{X}

We now analyse the case when \mathcal{X} is a compact and convex subset of $[0, r]^d$. For this, similar to the discrete case we will define two partitionings of the space \mathcal{X} . Let $\mathcal{J}_\eta^{(m)}$ be as defined in the discrete case. The first partitioning, $\{\mathcal{X}^{(m)}\}_{m=1}^M$ is problem dependent and is exactly as defined before. The second $\{\mathcal{H}_\tau^{(m)}\}_{m=1}^M$ is defined for the γ used in the algorithm and for any given $\rho, \tau > 0$. We make the dependence on τ explicit and write $\mathcal{H}_\tau^{(1)} = \overline{\mathcal{J}}_{\max(\tau, \rho\gamma)}^{(1)}$. Then define,

$$\mathcal{H}_\tau^{(m)} = \overline{\mathcal{J}}_{\max(\tau, \rho\gamma)}^{(m)} \cap \left(\bigcap_{\ell=1}^{m-1} \mathcal{J}_{\max(\tau, \rho\gamma)}^{(\ell)} \right) \quad \text{for } 2 \leq m \leq M-1, \quad \mathcal{H}_\tau^{(M)} = \bigcap_{\ell=1}^{M-1} \mathcal{J}_{\max(\tau, \rho\gamma)}^{(\ell)}. \quad (12)$$

In addition, as before define the sets above $\mathcal{H}_\tau^{(m)}$ as $\widehat{\mathcal{H}}_\tau^{(m)} = \bigcup_{\ell=m+1}^M \mathcal{H}_\tau^{(\ell)}$ and the sets below $\mathcal{H}_\tau^{(m)}$ as $\check{\mathcal{H}}_\tau^{(m)} = \bigcup_{\ell=1}^{m-1} \mathcal{H}_\tau^{(\ell)}$. Finally, for any given $\alpha > 0$ we will also define $\mathcal{H}_{\tau, n}^{(m)} = \{x \in \mathcal{X} : B_2(x, r\sqrt{d}/n^{\frac{\alpha}{2d}}) \cap \mathcal{H}_\tau^{(m)} \neq \emptyset \wedge x \notin \widehat{\mathcal{H}}_\tau^{(m)}\}$ to be an n -dependence inflation of $\mathcal{H}_\tau^{(m)}$. Here, $B_2(x, \epsilon)$ is an L_2 ball of radius ϵ centred at x . The sets $\{\mathcal{H}_{\tau, n}^{(m)}\}_{m=1}^M$ depend on ρ, γ, τ, n and α . Notice that for any $\alpha > 0$, as $n \rightarrow \infty$, $\mathcal{H}_{\tau, n}^{(m)} \rightarrow \mathcal{H}_\tau^{(m)}$. In addition to the above, denote the ε covering number of a set $A \subset \mathcal{X}$ in the $\|\cdot\|_2$ metric by $\Omega_\varepsilon(A)$. Our main theorem is as follows.

Theorem 13. *Let $\mathcal{X} \subset [0, r]^d$ be compact and convex. Let $f^{(m)} \sim \mathcal{GP}(\mathbf{0}, \kappa) \forall m$, and satisfy assumptions **A2**, **A3**. Let κ satisfy Assumption 8 with some constants a, b . Pick $\delta \in (0, 1)$ and run MF-GP-UCB with*

$$\beta_t = 2 \log \left(\frac{M\pi^2 t^2}{2\delta} \right) + 4d \log(t) + \max \left\{ 0, 2d \log \left(b r d \log \left(\frac{6M a d}{\delta} \right) \right) \right\}.$$

For all $\alpha \in (0, 1)$, $\tau > 0$, $\rho > \rho_0 = \max\{2, 1 + \sqrt{(1 + 2/\alpha)/(1 + d)}\}$ and sufficiently large Λ , we have $R(\Lambda) \in \mathcal{O} \left(\sum_{m=1}^M \lambda^{(m)} \sqrt{n_\Lambda \beta_{n_\Lambda} \Psi_{n_\Lambda}(\mathcal{H}_{\tau, n_\Lambda}^{(m)})} + \frac{\text{diam}(\widehat{\mathcal{H}}_\tau^{(m)})^d \text{polylog}(n_\Lambda)}{\text{poly}(\gamma)} \right)$. Here, $n_\Lambda = \lfloor \Lambda / \lambda^{(M)} \rfloor$ as before.

Precisely, there exists Λ_0 such that for all $\Lambda \geq \Lambda_0$, with probability $> 1 - \delta$ we have,

$$\begin{aligned} R(\Lambda) \leq & 2B\lambda^{(M)} + \lambda^{(M)} \left[\sqrt{2C_1 M n_\Lambda^\alpha \Psi_{2M n_\Lambda^\alpha}(\check{\mathcal{H}}^{(M)})} + \sqrt{2C_1 n_\Lambda \Psi_{2n_\Lambda}(\mathcal{H}_{\tau, n_\Lambda}^{(M)})} + \frac{\pi^2}{6} \right] \\ & + 2B \sum_{m=1}^{M-1} \lambda^{(m)} \left[(m-1)(2n_\Lambda^\alpha) + \frac{1}{\tau} \left(\sqrt{2C_1 n_\Lambda \beta_{2n_\Lambda} \Psi_{2n_\Lambda}(\mathcal{H}_{\tau, n_\Lambda}^{(m)})} + \frac{\pi^2}{6} \right) + \Omega_{\varepsilon_n}(\widehat{\mathcal{H}}_\tau^{(m)}) \left(\frac{2\eta^2}{\gamma^2} \beta_n + 1 \right) \right], \end{aligned}$$

where $C_1 = 8/\log(1 + \eta^2)$. For the SE kernel $\varepsilon_n = \frac{\gamma}{\sqrt{8C_{SE}\beta_n}}$, and therefore $\Omega_{\varepsilon_n}(\widehat{\mathcal{H}}_\tau^{(m)}) \in \mathcal{O} \left(\frac{\text{diam}(\widehat{\mathcal{H}}_\tau^{(m)})^d (\log(n))^{d/2}}{\gamma^d} \right)$. For the Matérn kernel $\varepsilon_n = \frac{\gamma^2}{8C_{Mat}\beta_n}$ and therefore $\Omega_{\varepsilon_n}(\widehat{\mathcal{H}}_\tau^{(m)}) \in \mathcal{O} \left(\frac{\text{diam}(\widehat{\mathcal{H}}_\tau^{(m)})^d (\log(n))^d}{\gamma^{2d}} \right)$. C_{SE}, C_{Mat} are kernel dependent constants. As $\Lambda \rightarrow \infty$, $n_\Lambda \rightarrow \infty$ and hence $\mathcal{H}_{\tau, n_\Lambda}^{(m)} \rightarrow \mathcal{H}_\tau^{(m)}$ for all $m \in \{1, \dots, M\}$ and $\alpha \in (0, 1)$.

Synopsis: Ignoring the common terms, constants and n_Λ^α terms, the regret for GP-UCB is $\lambda^{(M)} \sqrt{n_\Lambda \Psi_{n_\Lambda}(\mathcal{X})}$ whereas for MF-GP-UCB it is $\sum_m \lambda^{(m)} \sqrt{n_\Lambda \Psi_{n_\Lambda}(\mathcal{H}_{\tau, n}^{(m)})}$. In problems where $\text{vol}(\mathcal{H}_{\tau, n}^{(M)}) \ll \text{vol}(\mathcal{H}_{\tau, n}^{(M-1)}) \ll \dots \ll \text{vol}(\mathcal{H}_{\tau, n}^{(1)}) = \text{vol}(\mathcal{X})$, MF-GP-UCB achieves better regret than GP-UCB since $\lambda^{(m)} \ll \lambda^{(m+1)}$. The $\lambda^{(m)} \sqrt{n_\Lambda^\alpha \Psi_{n_\Lambda^\alpha}(\mathcal{H}_{\tau, n}^{(m)})}$ terms can be made arbitrarily small by picking large enough ρ , provided $\mathcal{H}_{\tau, n}^{(m)}$ is still small relative to \mathcal{X} . On the other hand the $\text{diam}(\widehat{\mathcal{H}}_\tau^{(m)}) \text{polylog}(n_\Lambda) / \text{poly}(\gamma)$ terms could be big if γ is too small. MF-GP-UCB requires that γ will be chosen large enough so that the above term remains small relative to $\sqrt{n_\Lambda \beta_{n_\Lambda} \Psi_{n_\Lambda}(\mathcal{H}_\tau^{(m)})}$ which is not too restrictive since we expect $\widehat{\mathcal{H}}_\tau^{(m)}$ to be much smaller than $\mathcal{H}_\tau^{(m)}$. While this bound illustrates the trade-off on the choice of γ – for small γ , the $1/\text{poly}(\gamma)$ term increases but the set sizes $\mathcal{H}^{(m)}$ decreases – it does not suggest a way to choose γ in practice since there are several problem dependent terms. Our heuristics for setting γ seemed to work well in practice (see Section 5).

Proof of Theorem 13. Once again, we will study MF-GP-UCB after n time steps regardless of the queried fidelities and bound \hat{R}_n . Then we will bound the number of plays N within capital Λ . For the analysis, at time n we will consider




Figure 9. Illustration of the sets $\{\mathcal{F}_n^{(\ell)}\}_{\ell=1}^{m-1}$ with respect to $\mathcal{H}_\tau^{(m)}$. The grid represents a $r\sqrt{d}/n^{\alpha/2d}$ covering of \mathcal{X} . The yellow region is $\widehat{\mathcal{H}}_\tau^{(m)}$. The area enclosed by the solid red line (excluding $\widehat{\mathcal{H}}_\tau^{(m)}$) is $\mathcal{H}_\tau^{(m)}$. $\mathcal{H}_{\tau,n}^{(m)}$, shown by a dashed red line, is obtained by inflating $\mathcal{H}_\tau^{(m)}$ by $r\sqrt{d}/n^{\alpha/2d}$. The grey shaded region represents $\bigcup_{\ell=1}^{m-1} \mathcal{F}_n^{(\ell)}$. By our definition, $\bigcup_{\ell=1}^{m-1} \mathcal{F}_n^{(\ell)}$ contains the cells which are entirely outside $\mathcal{H}_\tau^{(m)}$. However, the inflation $\mathcal{H}_{\tau,n}^{(m)}$ is such that $\widehat{\mathcal{H}}_\tau^{(m)} \cup \mathcal{H}_{\tau,n}^{(m)} \cup \bigcup_{\ell=1}^{m-1} \mathcal{F}_n^{(\ell)} = \mathcal{X}$. As $n \rightarrow \infty$, $\mathcal{H}_{\tau,n}^{(m)} \rightarrow \mathcal{H}_\tau^{(m)}$.

a $\frac{r\sqrt{d}}{2n^{\frac{\alpha}{2d}}}$ -covering of the space \mathcal{X} of size $n^{\frac{\alpha}{2}}$. For instance, if $\mathcal{X} = [0, r]^d$ a sufficient discretisation would be an equally spaced grid having $n^{\alpha/2d}$ points per side. Let $\{a_{i,n}\}_{i=1}^{n^{\frac{\alpha}{2}}}$ be the points in the covering, $F_n = \{A_{i,n}\}_{i=1}^{n^{\frac{\alpha}{2}}}$ be the cells in the covering, i.e. $A_{i,n}$ is the set of points which are closest to $a_{i,n}$ in the covering. Next we define another partitioning of the space similar in spirit to (12) using this partitioning. First let $F_n^{(1)} = \{A_{i,n} \in F_n : A_{i,n} \subset \mathcal{J}_{\max(\tau, \rho\gamma)}^{(1)}\}$. Next,

$$F_n^{(m)} = \left\{ A_{i,n} \in F_n : A_{i,n} \subset \overline{\mathcal{J}}_{\max(\tau, \rho\gamma)}^{(m)} \wedge A_{i,n} \notin \bigcup_{\ell=1}^{m-1} F_n^{(\ell)} \right\} \quad \text{for } 2 \leq m \leq M-1. \quad (13)$$

Note that $F_n^{(m)} \subset F_n$. We define the following *disjoint* subsets $\{\mathcal{F}_n^{(m)}\}_{m=1}^{M-1}$ of \mathcal{X} via $\mathcal{F}_n^{(m)} = \bigcup_{A_{i,n} \in F_n^{(m)}} A_{i,n}$. We have illustrated $\bigcup_{\ell=1}^{m-1} \mathcal{F}_n^{(\ell)}$ with respect to $\mathcal{H}_\tau^{(m)}$ in Figure 9. By noting that $\mathcal{H}_{\tau,n}^{(1)} = \mathcal{H}^{(1)}$ we make the following observation,

$$T_n^{(m)}(\mathcal{X}) \leq \sum_{\ell=1}^{m-1} T_n^{(m)}(\mathcal{F}_n^{(\ell)}) + T_n^{(m)}(\mathcal{H}_{\tau,n}^{(m)}) + T_n^{(m)}(\widehat{\mathcal{H}}_\tau^{(m)}). \quad (14)$$

This follows by noting that $\widehat{\mathcal{H}}_\tau^{(m)} \cup \mathcal{H}_{\tau,n}^{(m)} \subset \bigcup_{\ell=1}^{m-1} \mathcal{F}_n^{(\ell)}$ (See Fig. 9). To control \tilde{R}_n we will bound control each of these terms individually. First we focus on $\widehat{\mathcal{H}}_\tau^{(m)}$ for which we use the following lemma. The proof is given in Section C.2.1.

Lemma 14. *Let $f \sim \mathcal{GP}(\mathbf{0}, \kappa)$, $f : \mathcal{X} \rightarrow \mathbb{R}$ and we observe $y = f(x) + \epsilon$ where $\epsilon \sim \mathcal{N}(0, \eta^2)$. Let $A \subset \mathcal{X}$ such that its L_2 diameter $\text{diam}(A) \leq D$. Say we have n queries $(\mathbf{x}_t)_{t=1}^n$ of which s points are in A . Then the posterior variance of the GP, $\kappa'(x, x)$ at any $x \in A$ satisfies*

$$\kappa'(x, x) \leq \begin{cases} C_{SE}D^2 + \frac{\eta^2}{s} & \text{if } \kappa \text{ is the SE kernel,} \\ C_{Mat}D + \frac{\eta^2}{s} & \text{if } \kappa \text{ is the Matérn kernel,} \end{cases}$$

for appropriate constants C_{SE}, C_{Mat} .

First consider the SE kernel. At time t consider any $\varepsilon_n = \frac{\gamma}{\sqrt{8C_{SE}\beta_n}}$ covering $(B_i)_{i=1}^{\varepsilon_n}$ of $\widehat{\mathcal{H}}_\tau^{(m)}$. The number of queries inside any B_i of this covering at time n will be at most $\frac{2\eta^2}{\gamma^2}\beta_n + 1$. To see this, assume we have already queried $2\eta^2/\gamma^2 + 1$ times inside B_i at time $t \leq n$. By Lemma 14 the maximum variance in A_i can be bounded by

$$\max_{x \in A_i} \kappa_{t-1}^{(m)}(x, x) \leq C_{SE}(2\varepsilon_n)^2 + \frac{\eta^2}{T_t^{(m)}(A_i)} < \frac{\gamma^2}{\beta_n}.$$

Therefore, $\beta_t^{1/2}\sigma_{t-1}^{(m)}(x) \leq \beta_n^{1/2}\sigma_{t-1}^{(m)}(x) < \gamma$ and we will not query inside A_i until time n . Therefore, the number of m^{th} fidelity queries is bounded by $\Omega_{\varepsilon_n}(\widehat{\mathcal{H}}_\tau^{(m)}) \left(\frac{2\eta^2}{\gamma^2}\beta_n + 1 \right)$. The proof for the Matérn kernel follows similarly using $\varepsilon_n = \frac{\gamma^2}{8C_{Mat}\beta_n}$. Next, we bound $T_n^{(m)}(\mathcal{H}_{\tau,n}^{(m)})$ for which we will use the following Lemma. The proof is given in Section C.2.2.

Lemma 15. *For β_t as given in Theorem 13, we have the following with probability $> 1 - 5\delta/6$.*

$$\forall m \in \{1, \dots, M\}, \quad \forall t \geq 1, \quad \Delta^{(m)}(\mathbf{x}_t) = f_* - f^{(m)}(\mathbf{x}_t) \leq 2\beta_t\sigma_{t-1}^{(m)}(\mathbf{x}_t) + 1/t^2.$$

First, we will analyse the quantity $\tilde{R}_n^{(m)} = \sum_{\substack{t: \mathbf{m}_t = m \\ \mathbf{x}_t \in \mathcal{H}_{\tau, n}^{(m)}}} \Delta^{(m)}(\mathbf{x}_t)$ for $m < M$. Lemma 15 gives us $\tilde{R}_n^{(m)} \leq 2\beta_n^{1/2} \sum \sigma_{t-1}^{(m)}(\mathbf{x}_t) + \pi^2/6$. Then, using Lemma 9 and Jensen's inequality we have,

$$\left(\tilde{R}_n^{(m)} - \frac{\pi^2}{6} \right)^2 \leq 4\beta_t T_n^{(m)}(\mathcal{H}_{\tau, n}^{(m)}) \sum_{\substack{t: \mathbf{m}_t = m \\ \mathbf{x}_t \in \mathcal{H}_{\tau, n}^{(m)}}} (\sigma_{t-1}^{(m)})^2(\mathbf{x}_t) \leq C_1 \beta_t T_n^{(m)}(\mathcal{H}_{\tau, n}^{(m)}) \Psi_{T_n^{(m)}(\mathcal{H}_{\tau, n}^{(m)})}(\mathcal{H}_{\tau, n}^{(m)}). \quad (15)$$

We therefore have, $\tilde{R}_n^{(m)} \leq \sqrt{C_1 n \beta_n \Psi_n(\mathcal{H}_{\tau, n}^{(m)})} + \pi^2/6$ since trivially $T_n^{(m)}(\mathcal{H}_{\tau, n}^{(m)}) < n$. However, since $\Delta^{(m)}(x) > \tau$ for $x \in \mathcal{H}_{\tau, n}^{(m)}$ we have $T_n^{(m)}(\mathcal{H}_{\tau, n}^{(m)}) < \frac{1}{\tau} \left(\sqrt{C_1 n \beta_n \Psi_n(\mathcal{H}_{\tau, n}^{(m)})} + \pi^2/6 \right)$.

Remark 16. Since $\Psi_n(\cdot)$ is typically sublinear in n , it is natural to ask if we can recursively apply this to obtain a tighter bound on $T_n^{(m)}(\mathcal{H}_{\tau, n}^{(m)})$. For instance, since $\Psi_n(\cdot)$ is polylog(n) for the SE kernel (Srinivas et al. (2010), Theorem 5) by repeating the argument above once we get, $T_n^{(m)}(\mathcal{H}_{\tau, n}^{(m)}) \in \mathcal{O}\left(\frac{1}{\tau^{3/2}} \sqrt{C_1 n^{1/2} \text{polylog}(n) \beta_n \Psi_{\tau^{-3/2} n^{1/2} \text{polylog}(n)}}(\mathcal{H}_{\tau, n}^{(m)})\right)$. However, while this improves the dependence on n it worsens the dependence on τ . In fact, using a discretisation argument similar to that in Lemma 17 and the variance bound in Lemma 14, a polylog(n)/poly(τ) bound can be shown, with the poly(τ) term being τ^{d+2} for the SE kernel and τ^{2d+2} for the Matérn kernel.

Finally, to control the first term in (14), we will bound $T_n^{(>m)}(\mathcal{F}_n^{(m)})$. To that end we provide the following Lemma. The proof is given in Section C.2.3.

Lemma 17. Consider any $A_{i, n} \in F_n^{(m)}$ where $F_n^{(m)}$ is as defined in (13) for any $\alpha \in (0, 1)$. Let ρ, β_t be as given in Theorem 13, Then for all $u \geq \max\{3, (2(\rho - \rho_0)\eta)^{-2/3}\}$ we have,

$$\mathbb{P}(T_n^{(>m)}(A_{i, n}) > u) \leq \frac{\delta}{\pi^2} \cdot \frac{1}{u^{1+4/\alpha}}$$

We will use the above result with $u = n^{\alpha/2}$. Applying the union bound we have,

$$\begin{aligned} \mathbb{P}\left(\forall m \in \{1, \dots, M\}, T_n^{(>m)}(\mathcal{F}_n^{(m)}) > |F_n^{(m)}| n^{\alpha/2}\right) &\leq \sum_{m=1}^M \mathbb{P}\left(T_n^{(>m)}(\mathcal{F}_n^{(m)}) > |F_n^{(m)}| n^{\alpha/2}\right) \\ &\leq \sum_{m=1}^M \sum_{A_{i, n} \in F_n^{(m)}} \mathbb{P}\left(T_n^{(>m)}(A_{i, n}) > n^{\alpha/2}\right) \leq \sum_{m=1}^M |F_n^{(m)}| \frac{\delta}{\pi^2} \frac{1}{n^{2+\alpha/2}} \leq |F_n| \frac{\delta}{\pi^2} \frac{1}{n^{2+\alpha/2}} = \frac{\delta}{\pi^2} \frac{1}{n^2} \end{aligned}$$

Applying the union bound once again, we have $T_n^{(>m)}(\mathcal{F}_n^{(m)}) \leq n^\alpha$ for all m and all $n \geq \max\{3, (2(\rho - \rho_0)\eta)^{2/3}\}^{2/\alpha}$ with probability $> 1 - \delta/6$. Henceforth, all statements we make will make use of the results in Lemmas 14, 15 and 17 and will hold with probability $> 1 - \delta$.

First using equation (14) and noting $T_n^{(m)}(\mathcal{F}_n^{(\ell)}) \leq T_n^{(>\ell)}(\mathcal{F}_n^{(\ell)})$ for $\ell < m$ we bound $T_n^{(m)}(\mathcal{X})$ for $m < M$.

$$T_n^{(m)}(\mathcal{X}) \leq (m-1)n^\alpha + \frac{1}{\tau} \left(\sqrt{C_1 n \beta_n \Psi_n(\mathcal{H}_{\tau, n}^{(m)})} + \frac{\pi^2}{6} \right) + \Omega_{\varepsilon_n}(\hat{\mathcal{H}}^{(m)}) \left(\frac{2\eta^2}{\gamma^2} \beta_n + 1 \right).$$

Using this bound we can control $\tilde{R}_{n,1}$ in (5). To bound $\tilde{R}_{n,2}$ and $\tilde{R}_{n,3}$ we set $\mathcal{Z} = \mathcal{H}_{\tau, n_\Lambda}^{(M)}$ and use Lemma 15 noting that when $\mathbf{m}_t = M$, $r_t = \Delta^{(M)}(\mathbf{x}_t)$. Using similar calculations to (15) and as $T_n^{(M)}(\mathcal{H}_{\tau, n}^{(M)}) \leq n$, we have $\tilde{R}_{n,2} \leq \sqrt{C_1 n \beta_n \Psi_n(\mathcal{H}_{\tau, n}^{(M)})} + \sum_{\mathbf{x}_t \in \mathcal{Z}} 1/t^2$. Next, using Lemma 17 and observing $\overline{\mathcal{Z}} = \overline{\mathcal{H}_{\tau, n}^{(M)}} \subset \bigcup_{\ell=1}^{M-1} \mathcal{F}_n^{(m)} \subset \tilde{\mathcal{H}}^{(M)}$, we have,

$$\tilde{R}_{n,3} = \sum_{\substack{t: \mathbf{m}_t = M \\ \mathbf{x}_t \in \overline{\mathcal{Z}}}} (f_\star - f^{(M)}(\mathbf{x}_t)) \leq \sum_{\substack{t: \mathbf{m}_t = M \\ \mathbf{x}_t \in \bigcup_{\ell=1}^{M-1} \mathcal{F}_n^{(m)}}} 2\beta_t^{1/2} \sigma_{t-1}^{(m)}(\mathbf{x}_t) + \sum_{\mathbf{x}_t \in \overline{\mathcal{Z}}} \frac{1}{t^2} \leq \sqrt{C_1 M n^\alpha \beta_n \Psi_{M n^\alpha}(\tilde{\mathcal{H}}^{(M)})} + \sum_{\mathbf{x}_t \in \overline{\mathcal{Z}}} \frac{1}{t^2}.$$

Plugging these bounds back into (5), we obtain a bound on the regret similar to the one given in the theorem except with n replaced by $2n_\Lambda$. The last step in the proof will be to show that for sufficiently large Λ , $N \leq 2n_\Lambda$ which will complete the

proof. For this we turn back to our bounds for $T_n^{(m)}(\mathcal{X})$, $m < M$. Using a similar reasoning to (9), we can show that the following term upper bounds the number of queries at fidelities less than M ,

$$(M-1)n^\alpha + \sum_{m=1}^{M-1} \frac{1}{\tau} \left(\sqrt{2C_1 n_\Lambda \beta_{2n_\Lambda} \Psi_{2n_\Lambda}(\mathcal{H}_{\tau, n_\Lambda}^{(m)})} + \frac{\pi^2}{6} \right) + \sum_{m=1}^{M-1} \Omega_{\varepsilon_n}(\widehat{\mathcal{H}}^{(m)}) \left(\frac{2\eta^2}{\gamma^2} \beta_n + 1 \right).$$

Assume n_0 is large enough so that $n_0 \geq \max\{3, (2(\rho - \rho_0)\eta)^{-2/3}\}^{2/\alpha}$ and for all $n \geq n_0$, $n/2$ is larger than the above upper bound. We can find such an n_0 since the bound is $o(n)$. Therefore, for all $n \geq n_0$, $T_n^{(M)}(\mathcal{X}) > n/2$. Following a similar argument to the discrete case, $N \leq 2\Lambda/\lambda^{(M)}$ with probability $> 1 - \delta$ if $\Lambda \geq \Lambda_0 = \lambda^{(M)}(n_0 + 1)$. The theorem follows with the observation $N \geq n_\Lambda \implies \mathcal{H}_{\tau, N}^{(m)} \subset \mathcal{H}_{\tau, n_\Lambda}^{(m)} \implies \Psi_N(\mathcal{H}_{\tau, N}^{(m)}) \leq \Psi_N(\mathcal{H}_{\tau, n_\Lambda}^{(m)}) \leq \Psi_{2n_\Lambda}(\mathcal{H}_{\tau, n_\Lambda}^{(m)})$. ■

C.2.1. PROOF OF LEMMA 14

Since the posterior variance only decreases with more observations, we can upper bound $\kappa'(x, x)$ for any $x \in A$ by considering its posterior variance with only the s observations in A . Next the maximum variance within A occurs if we pick 2 points x_1, x_2 that are distance D apart and have all observations at x_1 ; then x_2 has the highest posterior variance. Therefore, we will bound $\kappa'(x, x)$ for any $x \in A$ with $\kappa(x_2, x_2)$ in the above scenario. Let $\kappa_0 = \kappa(x, x)$ and $\kappa(x, x') = \kappa_0 \phi(\|x - x'\|_2)$, where $\phi(\cdot) \leq 1$ depends on the kernel. Denote the gram matrix in the scenario described above by $\Delta = \kappa_0 \mathbf{1}\mathbf{1}^\top + \eta^2 I$. Then using the Sherman-Morrison formula, the posterior variance (1) can be bounded via,

$$\begin{aligned} \kappa'(x, x) &\leq \kappa'(x_2, x_2) = \kappa(x_2, x_2) - [\kappa(x_1, x_2)\mathbf{1}]^\top \Delta^{-1} [\kappa(x_1, x_2)\mathbf{1}] = \kappa_0 - \kappa_0^2 \phi^2(D) \mathbf{1}^\top [\kappa_0 \mathbf{1}\mathbf{1}^\top + \eta^2 I]^{-1} \mathbf{1} \\ &= \kappa_0 - \kappa_0 \phi^2(D) \mathbf{1}^\top \left[\frac{\kappa_0}{\eta^2} I - \frac{\left(\frac{\kappa_0}{\eta^2}\right)^2 \mathbf{1}\mathbf{1}^\top}{1 + \frac{\kappa_0}{\eta^2} s} \right] \mathbf{1} = \kappa_0 - \kappa_0 \phi^2(D) \left(\frac{\kappa_0}{\eta^2} s - \frac{\left(\frac{\kappa_0}{\eta^2}\right)^2 s^2}{1 + \frac{\kappa_0}{\eta^2} s} \right) \\ &= \kappa_0 - \kappa_0 \phi^2(D) \frac{s}{\frac{\eta^2}{\kappa_0} + s} = \frac{1}{1 + \frac{\eta^2}{\kappa_0 s}} \left(\kappa_0 - \kappa_0 \phi^2(D) + \frac{\eta^2}{s} \right) \\ &\leq \kappa_0 (1 - \phi^2(D)) + \frac{\eta^2}{s}. \end{aligned}$$

For the SE kernel $\phi^2(D) = \exp\left(\frac{-D^2}{2h^2}\right) = \exp\left(\frac{-D^2}{h^2}\right) \leq 1 - \frac{D^2}{h^2}$. Plugging this into the bound above retrieves the first result with $C_{SE} = \kappa_0/h^2$. For the Matérn kernel we use a Lipschitz constant L_{Mat} of ϕ . Then $1 - \phi^2(D) = (1 - \phi(D))(1 + \phi(D)) \leq 2(\phi(0) - \phi(D)) \leq 2L_{Mat}D$. We get the second result with $C_{Mat} = 2\kappa_0 L_{Mat}$. Since the SE kernel decays fast, we get a stronger result on its posterior variance which translates to a better bound in our theorems. ■

C.2.2. PROOF OF LEMMA 15

The first part of the proof mimics the arguments in Lemmas 5.6, 5.7 of Srinivas et al. (2010). By assumption 8 and the union bound we can show,

$$\mathbb{P}\left(\forall m \in \{1, \dots, M\}, \forall i \in \{1, \dots, d\}, \forall x \in \mathcal{X}, \left| \frac{\partial f^{(m)}(x)}{\partial x_i} \right| < b \log\left(\frac{6Mad}{\delta}\right)\right) \geq 1 - \frac{\delta}{6}.$$

Now we construct a discretisation F_t of \mathcal{X} of size $(\nu_t)^d$ such that we have for all $x \in \mathcal{X}$, $\|x - [x]_t\|_1 \leq rd/\nu_t$. Here $[x]_t$ is the closest point to x in the discretisation. (Note that this is different from the discretisation appearing in Theorem 13 even though we have used the same notation). By choosing $\nu_t = t^2 brd \sqrt{6Mad/\delta}$ and using the above we have

$$\forall x \in \mathcal{X}, \quad |f^{(m)}(x) - f^{(m)}([x]_t)| \leq b \log(6Mad/\delta) \|x - [x]_t\|_1 \leq 1/t^2 \quad (16)$$

for all $f^{(m)}$'s with probability $> 1 - \delta/6$.

Noting that $\beta_t \geq 2 \log(M|F_t|\pi^2 t^2/2\delta)$ for the given choice of ν_t we have the following with probability $> 1 - \delta/3$.

$$\forall t \geq 1, \forall m \in \{1, \dots, M\}, \forall a \in F_t, \quad |f^{(m)}(a) - \mu_{t-1}^{(m)}(a)| \leq \beta_t^{1/2} \sigma_{t-1}^{(m)}(a). \quad (17)$$

The proof mimics that of Lemma 11 using the same conditioning argument. However, instead of a fixed set over all t , we change the set at which we have confidence based on the discretisation. Similarly we can show that with probability $> 1 - \delta/3$ we also have confidence on the decisions \mathbf{x}_t at all time steps. Precisely,

$$\forall t \geq 1, \forall m \in \{1, \dots, M\}, \quad |f^{(m)}(\mathbf{x}_t) - \mu_{t-1}^{(m)}(\mathbf{x}_t)| \leq \beta_t^{1/2} \sigma_{t-1}^{(m)}(\mathbf{x}_t). \quad (18)$$

Using (16), (17) and (18) the following statements hold with probability $> 1 - 5\delta/6$. First we can upper bound f_* by,

$$f_* \leq f^{(m)}(x_*) + \zeta^{(m)} \leq f^{(m)}([x_*]_t) + \zeta^{(m)} + \frac{1}{t^2} \leq \varphi_t^{(m)}([x_*]_t) + \frac{1}{t^2}. \quad (19)$$

Since the above holds for all m , we have $f_* \leq \varphi_t([x_*]_t) + 1/t^2$. Now, using similar calculations as (7) we bound $\Delta^{(m)}(\mathbf{x}_t)$.

$$\begin{aligned} \Delta^{(m)}(\mathbf{x}_t) &= f_* - f^{(m)}(\mathbf{x}_t) - \zeta^{(m)} \leq \varphi_t([x_*]_t) + \frac{1}{t^2} - f^{(m)}(\mathbf{x}_t) - \zeta^{(m)} \leq \varphi_t(\mathbf{x}_t) - f^{(m)}(\mathbf{x}_t) - \zeta^{(m)} + \frac{1}{t^2} \\ &\leq \varphi_t^{(m)}(\mathbf{x}_t) - \mu_{t-1}^{(M)}(\mathbf{x}_t) + \beta_t^{1/2} \sigma_{t-1}^{(M)}(\mathbf{x}_t) - \zeta^{(m)} + \frac{1}{t^2} \leq 2\beta_t^{1/2} \sigma_{t-1}^{(M)}(\mathbf{x}_t) + \frac{1}{t^2}. \end{aligned}$$

C.2.3. PROOF OF LEMMA 17

First, we will invoke the same discretisation used in the proof of Lemma 15 via which we have $\varphi_t([x_*]_t) \geq f_* - 1/t^2$ (19). (Therefore, Lemma 17 holds only with probability $> 1 - \delta/6$, but this event has already been accounted for in Lemma 15.) Let $b_{i,n,t} = \operatorname{argmax}_{x \in A_{i,n}} \varphi_t(x)$ be the maximiser of the upper confidence bound in $A_{i,n}$ at time t . Now using the relaxation $\mathbf{x}_t \in A_{i,n} \implies \varphi_t(b_{i,n,t}) > \varphi_t([x_*]_t) \implies \varphi_t^{(m)}(b_{i,n,t}) > f_* - 1/t^2$ and proceeding,

$$\begin{aligned} \mathbb{P}(T_n^{(>m)}(A_{i,n}) > u) &\leq \mathbb{P}(\exists t : u+1 \leq t \leq n, \varphi_t^{(m)}(b_{i,n,t}) > f_* - 1/t^2 \quad \wedge \quad \beta_t^{1/2} \sigma_{t-1}^{(m)}(b_{i,n,t}) < \gamma) \\ &\leq \sum_{t=u+1}^n \mathbb{P}(\mu_{t-1}^{(m)}(b_{i,n,t}) - f^{(m)}(b_{i,n,t}) > \Delta^{(m)}(b_{i,n,t}) - \beta_t^{1/2} \sigma_{t-1}^{(m)}(b_{i,n,t}) - 1/t^2 \quad \wedge \quad \beta_t^{1/2} \sigma_{t-1}^{(m)}(b_{i,n,t}) < \gamma) \\ &\leq \sum_{t=u+1}^n \mathbb{P}(\mu_{t-1}^{(m)}(b_{i,n,t}) - f^{(m)}(b_{i,n,t}) > (\rho - 1)\beta_t^{1/2} \sigma_{t-1}^{(m)}(b_{i,n,t}) - 1/t^2) \\ &\leq \sum_{t=u+1}^n \mathbb{P}_{Z \sim \mathcal{N}(0,1)}\left(Z > (\rho_0 - 1)\beta_t^{1/2}\right) \leq \sum_{t=u+1}^n \frac{1}{2} \exp\left(-\frac{(\rho_0 - 1)^2}{2} \beta_t\right) \\ &\leq \frac{1}{2} \left(\frac{\delta}{M\pi^2}\right)^{(\rho_0 - 1)^2} \sum_{t=u+1}^n t^{-(\rho_0 - 1)^2(2+2d)} \leq \frac{\delta}{M\pi^2} u^{-(\rho_0 - 1)^2(2+2d)+1} \leq \frac{\delta}{\pi^2} \frac{1}{u^{1+4/\alpha}}. \end{aligned} \quad (20)$$

In the second step we have rearranged the terms and used the definition of $\Delta^{(m)}(x)$. In the third step, as $A_{i,n} \subset \mathcal{F}_{\max(\tau, \rho\gamma)}^{(m)}$, $\Delta^{(m)}(b_{i,n,t}) > \rho\gamma > \rho\beta_t^{1/2} \sigma_{t-1}^{(m)}(b_{i,n,t})$. In the fourth step we have used the following facts, $t > u \geq \max\{3, (2(\rho - \rho_0)\eta)^{-2/3}\}$, $M\pi^2/2\delta > 1$ and $\sigma_{t-1}^{(m)}(b_{i,n,t}) > \eta/\sqrt{t}$ to conclude,

$$(\rho - \rho_0) \frac{\eta\sqrt{4\log(t)}}{\sqrt{t}} > \frac{1}{t^2} \implies (\rho - \rho_0) \cdot \sqrt{2\log\left(\frac{M\pi^2 t^2}{2\delta}\right)} \cdot \frac{\eta}{\sqrt{t}} > \frac{1}{t^2} \implies (\rho - \rho_0)\beta_t^{1/2} \sigma_{t-1}^{(m)}(b_{i,n,t}) > \frac{1}{t^2}.$$

In the seventh step of (20) we have bound the sum by an integral and used $\rho_0 \geq 2$ twice. Finally, the last step follows by $\rho_0 \geq 1 + \sqrt{(1 + 2/\alpha)/(1 + d)}$ and noting $M \geq 1$. ■

D. Addendum to Experiments

D.1. Description of Synthetic Experiments

The following are the descriptions of the synthetic functions used. The first three functions and their approximations were taken from (Xiong et al., 2013).

Currin exponential function: The domain is $\mathcal{X} = [0, 1]^2$. The second and first fidelity functions are,

$$\begin{aligned} f^{(2)}(x) &= \left(1 - \exp\left(\frac{-1}{2x_2}\right)\right) \left(\frac{2300x_1^3 + 1900x_1^2 + 2092x_1 + 60}{100x_1^3 + 500x_1^2 + 4x_1 + 20}\right), \\ f^{(1)}(x) &= \frac{1}{4}f^{(2)}(x_1 + 0.05, x_2 + 0.05) + \frac{1}{4}f^{(2)}(x_1 + 0.05, \max(0, x_2 - 0.05)) + \\ &\quad \frac{1}{4}f^{(2)}(x_1 - 0.05, x_2 + 0.05) + \frac{1}{4}f^{(2)}(x_1 - 0.05, \max(0, x_2 - 0.05)). \end{aligned}$$

Park function: The domain is $\mathcal{X} = [0, 1]^4$. The second and first fidelity functions are,

$$\begin{aligned} f^{(2)}(x) &= \frac{x_1}{2} \left(\sqrt{1 + (x_2 + x_3^2) \frac{x_4}{x_1^2}} - 1 \right) + (x_1 + 3x_4) \exp(1 + \sin(x_3)), \\ f^{(1)}(x) &= \left(1 + \frac{\sin(x_1)}{10}\right) f^{(2)}(x) - 2x_1^2 + x_2^2 + x_3^2 + 0.5. \end{aligned}$$

Borehole function: The second and first fidelity functions are,

$$f^{(2)}(x) = \frac{2\pi x_3(x_4 - x_6)}{\log(x_2/x_1) \left(1 + \frac{2x_7x_3}{\log(x_2/x_1)x_1^2x_8} + \frac{x_3}{x_5}\right)}, \quad f^{(1)}(x) = \frac{5x_3(x_4 - x_6)}{\log(x_2/x_1) \left(1.5 + \frac{2x_7x_3}{\log(x_2/x_1)x_1^2x_8} + \frac{x_3}{x_5}\right)}.$$

The domain of the function is $[0.05, 0.15; 100, 50K; 63.07K, 115.6K; 990, 1110; 63.1, 116; 700, 820; 1120, 1680; 9855, 12045]$ but we first linear transform the variables to lie in $[0, 1]^8$.

Hartmann-3D function: The M^{th} fidelity function is $f^{(M)}(x) = \sum_{i=1}^4 \alpha_i \exp\left(-\sum_{j=1}^3 A_{ij}(x_j - P_{ij})^2\right)$ where $A, P \in \mathbb{R}^{4 \times 3}$ are fixed matrices given below and $\alpha = [1.0, 1.2, 3.0, 3.2]$. For the lower fidelities we use the same form except change α to $\alpha^{(m)} = \alpha + (M - m)\delta$ where $\delta = [0.01, -0.01, -0.1, 0.1]$ and $M = 3$. The domain is $\mathcal{X} = [0, 1]^3$.

$$A = \begin{bmatrix} 3 & 10 & 30 \\ 0.1 & 10 & 35 \\ 3 & 10 & 30 \\ 0.1 & 10 & 35 \end{bmatrix}, \quad P = 10^{-4} \times \begin{bmatrix} 3689 & 1170 & 2673 \\ 4699 & 4387 & 7470 \\ 1091 & 8732 & 5547 \\ 381 & 5743 & 8828 \end{bmatrix}$$

Hartmann-6D function: The 6-D Hartmann takes the same form as above except $A, P \in \mathbb{R}^{4 \times 6}$ are as given below. We use the same modification to obtain the lower fidelities using $M = 4$.

$$A = \begin{bmatrix} 10 & 3 & 17 & 3.5 & 1.7 & 8 \\ 0.05 & 10 & 17 & 0.1 & 8 & 14 \\ 3 & 3.5 & 1.7 & 10 & 17 & 8 \\ 17 & 8 & 0.05 & 10 & 0.1 & 14 \end{bmatrix}, \quad P = 10^{-4} \times \begin{bmatrix} 1312 & 1696 & 5569 & 124 & 8283 & 5886 \\ 2329 & 4135 & 8307 & 3736 & 1004 & 9991 \\ 2348 & 1451 & 3522 & 2883 & 3047 & 6650 \\ 4047 & 8828 & 8732 & 5743 & 1091 & 381 \end{bmatrix}$$

D.2. More Results on Synthetic Experiments

Figure 10 shows the simple regret $S(\Lambda)$ for the synthetic functions not presented in the main text.

It is natural to ask how MF-GP-UCB performs with bad approximations at lower fidelities. We found that our implementation with the heuristics suggested in Section 5 to be quite robust. We demonstrate this using the Currin exponential function, but using the negative of $f^{(2)}$ as the first fidelity approximation, i.e. $f^{(1)}(x) = -f^{(2)}(x)$. Figure 11 illustrates $f^{(1)}, f^{(2)}$ and gives the simple regret $S(\Lambda)$. Understandably, it loses to the single fidelity methods since the first fidelity queries are wasted and it spends some time at the second fidelity recovering from the bad approximation. However, it eventually is able to achieve low regret.

Finally, we present results on the cumulative regret for the synthetic functions in Figure 12.




Figure 10. The simple regret $S(\Lambda)$ against the spent capital Λ on the synthetic functions. The title states the function, its dimensionality, the number of fidelities and the costs we used for each fidelity in the experiment. All curves barring DiRect (which is a deterministic), were produced by averaging over 20 experiments. The error bars indicate one standard error.




Figure 11. (a) illustrates the functions used in the Bad Currin Exponential experiment where we took $f^{(1)} = -f^{(2)}$ and (b) shows the simple regret for this experiment. See caption under Fig. 10 for more details.

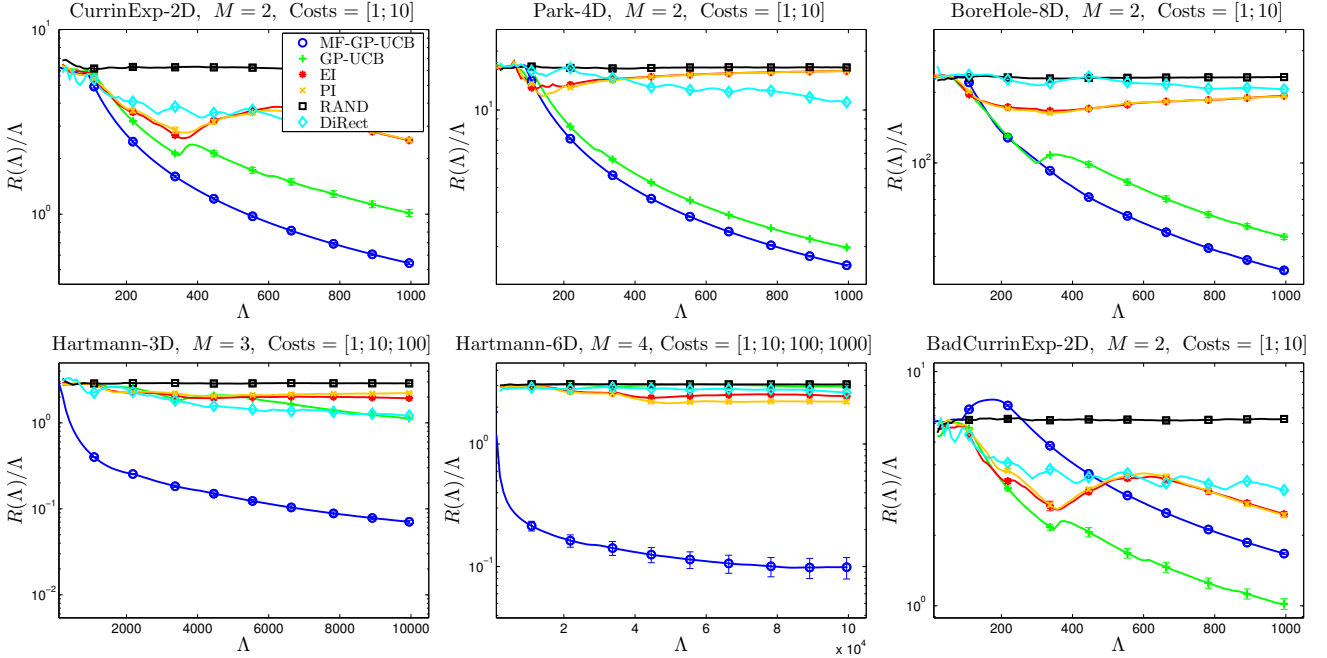


Figure 12. The cumulative regret $R(\Lambda)$ against the spent capital Λ on the synthetic functions. The title states the function, its dimensionality, the number of fidelities and the costs we used for each fidelity in the experiment. All curves barring DiRect (which is a deterministic), were produced by averaging over 20 experiments. The error bars indicate one standard error.